IPv6 Fundamentals

A Straightforward Approach to Understanding IPv6

Second Edition

ciscopress.com

Rick Graziani

IPv6 Fundamentals: A Straightforward Approach to Understanding IPv6

Second Edition

Rick Graziani

Cisco Press

800 East 96th Street

Indianapolis, IN 46240

IPv6 Fundamentals: A Straightforward Approach to Understanding IPv6, Second Edition

Rick Graziani

Copyright © 2017 Cisco Systems, Inc.

Cisco Press logo is a trademark of Cisco Systems, Inc.

Published by:
Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

Printed in the United States of America

1 17

Library of Congress Control Number: 2017931983

ISBN-13: 978-1-58714-477-6

ISBN-10: 1-58714-477-8

Warning and Disclaimer

This book is designed to provide information about IPv6 (Internet Protocol version 6). Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an "as is" basis. The author, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc. cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Special Sales

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Editor-in-Chief: Mark Taub

Product Line Manager: Brett Bartow

Business Operation Manager, Cisco Press: Ronald Fligge

Executive Editor: Mary Beth Ray

Managing Editor: Sandra Schroeder

Development Editor: Ellie Bru

Project Editor: Mandie Frank

Copy Editor: Kitty Wilson

Technical Editors: Jim Bailey, Tim Martin

Editorial Assistant: Vanessa Evans

Cover Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Cheryl Lenser

Proofreader: Larry Sulky

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

About the Author

Rick Graziani has been an instructor of computer networking and computer science courses at Cabrillo College in Aptos, California since 1994. Rick also teaches networking courses in the Computer Engineering department at the University of California, Santa Cruz and is on the Curriculum Engineering team for Cisco Networking Academy. Prior to teaching, he worked in the information technology field for Santa Cruz Operation, Tandem Computers, Lockheed Missiles and Space Company, and served five years in the U.S. Coast Guard. When he is not working, he is most likely surfing at one of his favorite Santa Cruz surf breaks or hanging out with his dog, Luigi. You are welcome to use his instructional resources on his Cabrillo College website, www.cabrillo.edu/~rgraziani, for IPv6, CCNA, or CCNP information. You can email graziani@cabrillo.edu to obtain the username and password for all his materials.

Rick and Luigi

About the Technical Reviewers

Jim Bailey, CCIE No. 5275 (Routing & Switching; Service Provider) and CCDE No. 20090008, is a Solution Architect at Cisco Systems with more than 25 years of experience in networking. As part of the Cisco Advanced Services team, he works on the architecture, design, and implementation of networks for enterprise, service provider, and government customers. He has focused on IPv6 integration into those networks for more than 12 years. He has presented *IPv6 Planning, Deployment and Operational Considerations* at Cisco Live conferences. He has served as a technical reviewer for the *IPv6 for Enterprise Networks* and *IPv6 Design and Deployment LiveLessons* published by Cisco Press.

Tim Martin, CCIE No. 2020, is a dynamic presenter and a member of the Cisco Live Distinguished Speaker hall of fame. Tim frequently speaks at Cisco Live events in both the United States and Europe. Tim has been in the inter-networking industry for more than 34 years. Cisco Press recently published a title of his work *IPv6 Design & Deployment Live Lessons* (ISBN 9780134655512), a 6-hour video series that provides guidance on IPv6 enterprise design. In his current role at Cisco, he is a Solutions Architect focused on the U.S. public sector market. Tim achieved the distinction of a multi-protocol CCIE No. 2020 in June 1996 and has also attained the Gold Certified Engineer status from the IPv6 Forum. He has participated in numerous industry events related to IPv6 and contributes to the IETF IPv6 subcommittees. Tim is also a member of many different IPv6-related task forces, including the FEDv6TF, NAv6TF, TXv6TF, and RMv6TF.

Dedication

To brothers Frank and Mark. You are not only my brothers but you are my best friends. I love you both. Also, to all of my current and former students. I am humbled by the opportunity to teach such wonderful people. You make my job fun, and you are the reason I love to go to work every day.

Acknowledgments

First of all, I would like to thank my family their love and support. Family is the best.

I would like to thank all my friends and colleagues for their assistance. Mark Boolootian, Dave Barnett, and Jim Warner, thanks for the many years of discussing technologies and answering questions. We've drawn a lot of topologies on a lot of napkins over the years.

The technical editors Jim Bailey and Tim Martin at Cisco Systems deserve much more credit than the brief mention as technical reviewers for this book. They make me look a lot smarter than I am. They did an incredibly thorough job making sure that this book is as accurate and as up to date as possible. Their expertise and experience were invaluable in helping me author this book. They are both the unsung heroes of this project. Thank you for your dedication and your commitment.

I owe a great deal of gratitude to Gerlinde Brady, Mike Matera, and Rich Simms at Cabrillo College for their friendship and support. You made sure that the CS/CIS departments at Cabrillo College continued to run smoothly while I was engaged in the writing process. I feel very fortunate to work with all of you and all of our other friends in the CS/CIS department. Thank you David Hovey and Ahmad Allulu for lab support. Thank you Brad Smith, Patrick Mantey, J.J. Garcia-Luna Aceves, and Katia Obraczka for the privilege of teaching in the Computer Science and Engineering department at the University of California, Santa Cruz. Teaching students and working with faculty at UCSC have made this a much better book.

Writing this book has been one of many privileges I have received due to the honor of working with Dennis Frezzo, Jeremy Creech, Telethia Willis, and many others who work for the Cisco Networking Academy. Thank you all for the proud opportunity to be part of a program that has changed the lives of thousands of students around the world. More than colleagues, you are all friends, for which I am grateful.

Thank you, Pat Farley, for making sure I still get in my surf sessions. Now that the book is done, you'll see me in the lineup much more often. Thank you, Teri Graziani, for always being there and taking care of things while I was busy writing this book. I appreciate it more than you know.

Special thanks to Mary Beth Ray, executive editor for Cisco Press and a friend. Thank you for your patience and understanding through this long process. You always have that calm assurance and guidance.

Thank you, Ellie Bru at Cisco Press, for working with me on a daily basis—weekdays and weekends—editing, formatting, and orchestrating the entire process. You were a pleasure to work with, and I am very grateful for all the hard work you put into this book. To Mandie Frank, Dhayanidhi, Kitty Wilson, and Larry Sulky, thank you for making me look like a better writer than I really am. And to everyone else at Cisco Press, I am extremely grateful for everything you have done. I am constantly amazed at the level of cooperation and teamwork required to produce a technical book. I am very thankful for all your help.

Finally, thank you to all my friends in Rome and Frascati, Italy, where I spent many months writing this book. Thank you, Giuseppe Cinque, for your many years of friendship and for convincing me to spend my sabbatical in Italy writing this book. Thank you, Mama and Papa Cinque, for making me part of your Positano family. Thank you, Levi Adam, Fidele Lassandro, Antonio Brancaccia, Daniel and Andrea and everyone at Tusculum Sport Center, Fermate N'Attimo, everyone at Il Borgo Verde, the Molinari family at Antico Forno, and everyone at Etabli for your friendship. And a special thank you to Alice, Mauro, Loredana, Marco, and Timmy Chialastri, for making Rome a second home. Your kindness made the time I spent in Rome some of the most enjoyable months I have ever had. Thank you for opening up your home and your hearts to me. I will forever be grateful for everything you did for me.

Contents at a Glance

Introduction xxv

Part I Introduction to IPv6 1

Chapter 1 Introduction to IPv6 3

Chapter 2 IPv6 Primer 33

Chapter 3 Comparing IPv4 and IPv6 49

Part II IPv6 Addresses 89

Chapter 4 IPv6 Address Representation and Address Types 91

Chapter 5 Global Unicast Address 125

Chapter 6 Link-Local Unicast Address 167

Chapter 7 Multicast Addresses 193

Part III Dynamic IPv6 Addressing 225

Chapter 8 Basics of Dynamic Addressing in IPv6 227

Chapter 9 Stateless Address Autoconfiguration (SLAAC) 251

Chapter 10 Stateless DHCPv6 297

Chapter 11 Stateful DHCPv6 315

Part IV ICMPv6 and ICMPv6 Neighbor Discovery 345

Chapter 12 ICMPv6 347

Chapter 13 ICMPv6 Neighbor Discovery 373

Part V Routing IPv6 413

Chapter 14 IPv6 Routing Table and Static Routes 415

Chapter 15 EIGRP for IPv6 443

Chapter 16 OSPFv3 475

Part VI Implementing IPv6 515

Chapter 17 Deploying IPv6 in the Network 517

Appendixes

Appendix A Configuring NAT64 and IPv6 Tunnels 573

Appendix B IPv6 Command Quick Reference 601

Appendix C Answers to Review Questions 615

Index 631

Contents

Introduction xxv

Part I Introduction to IPv6 1

Chapter 1 Introduction to IPv6 3

IPv6 Is Here 3

Why Transition to IPv6? 5

IPv4 Address Depletion 6

Access to IPv6-Only Customers 6

Better Performance 6

Securing Your Current Network 7

IPv4 8

IPv4 Address Depletion 8

CIDR 11

NAT with Private Addresses 13

Problems with NAT 15

NAT is Not Security 16

NAT Example 17

What About IPv5? 19

The Fascinating History of IPv6 19

Some Background 20

IPv4 Address Exhaustion and the Need for More International Involvement 21

A Call for Proposals 22

A More IP Version of IPv6 23

IPv6: More Than Just Longer Addresses 24

IPv6 Myths 25

Transitioning to IPv6 26

Summary 28

Review Questions 28

References 29

Endnotes 29

RFCs 29

Websites 31

Chapter 2 IPv6 Primer 33

Hexadecimal Number System 34

IPv6 Address Types 37

Global Unicast Address (GUA) 37

Link-Local Unicast Address 37

Unspecified Address 38

Solicited-Node Multicast Address 38

Address Terminology 41

ICMPv6 Neighbor Discovery Protocol (NDP) 41

Neighbor Solicitation (NS) and Neighbor Advertisement (NA) Messages 42

Router Solicitation (RS) and Router Advertisement (RA) Messages 42

Dynamic Address Allocation 43

Summary 45

Review Questions 46

References 48

RFCs 48

Chapter 3 Comparing IPv4 and IPv6 49

Comparing the IPv4 and IPv6 Headers 49

The IPv4 and IPv6 Version Fields 51

IPv4 Internet Header Length (IHL) Field 51

IPv4 Type of Service (ToS) and IPv6 Traffic Class Fields 52

IPv6 Flow Label Field 54

IPv4 Total Length Field, IPv6 Payload Length Field, and IPv6 Jumbograms 54

IPv4 and IPv6 MTUs 56

IPv4 Fragmentation 57

IPv6 Fragmentation: IPv6 Source Only 58

IPv4 Protocol and IPv6 Next Header Fields 59

IPv4 Time to Live (TTL) and IPv6 Hop Limit Fields 62

Checksums: IPv4, TCP, and UDP 63

IPv4 and IPv6 Source Address and Destination Address Fields 65

IPv4 Options and Padding Fields, IPv6 Fixed Length 65

IPv6 over Ethernet 66

Packet Analysis Using Wireshark 66

Extension Headers 69
Hop-by-Hop Options Extension Header 72
Routing Extension Header 74
Fragment Extension Header 76
IPsec: AH and ESP Extension Headers 77
Transport and Tunnel Modes 78
Encapsulating Security Payload (ESP) Extension Header 79
Authentication Header (AH) Extension Header 81
Destination Options Extension Header 82
No Next Header 84
Comparing IPv4 and IPv6 at a Glance 84
Summary 86
Review Questions 86
References 86
RFCs 86
Websites 87

Part II IPv6 Addresses 89

Chapter 4 IPv6 Address Representation and Address Types 91
Representation of IPv6 Addresses 91
Rule 1: Omit Leading 0s 93
Rule 2: Omit All-0s Hextets 95
Combining Rule 1 and Rule 2 96
Prefix Length Notation 98
IPv6 Address Types 99
IPv6 Address Space 100
Unicast Addresses 103
Global Unicast Address 104
Link-Local Unicast Address 106
Loopback Addresses 109
Unspecified Addresses 109
Unique Local Addresses 110
ULA and NAT 111
L Flag and Global ID 112
Site-Local Addresses (Deprecated) 113
IPv4 Embedded Address 114
IPv4-Mapped IPv6 Addresses 114
IPv4-Compatible IPv6 Addresses (Deprecated) 115

Multicast Addresses 115
Well-Known Multicast Addresses 117
Solicited-Node Multicast Addresses 118
Anycast Addresses 118
Summary 119
Review Questions 121
References 122
Endnote 122
RFCs 122
Websites 123
Book 123

Chapter 5 Global Unicast Address 125
Structure of a Global Unicast Address 126
Global Routing Prefix 128
Subnet ID 129
Interface ID 129
Manual Configuration of a Global Unicast Address 130
Manual GUA Configuration for Cisco IOS 131
Manual GUA Configuration with EUI-64 for Cisco IOS 135
Manual GUA Configuration with IPv6 Unnumbered for Cisco IOS 137
Manual GUA Configuration for Windows, Linux, and Mac OS 137
Implementing Static Routing and Verifying Connectivity with Ping 141
Recognizing the Parts of a GUA Address and the 3–1–4 Rule 142
Examining Other Prefix Lengths 144
Subnetting IPv6 145
Extending the Subnet Prefix 148
Subnetting on a Nibble Boundary 149
Subnetting Within a Nibble 150
Subnetting /127 Point-to-Point Links 151
NDP Exhaustion Attack 151
/127 Subnetting on Point-to-Point Links 152
ipv6gen: An IPv6 Subnetting Tool 155
Prefix Allocation 156
Provider-Aggregatable (PA) and Provider-Independent (PI) Address Space 158
Provider-Aggregatable Address Space 158
Provider-Independent Address Space 159

General Prefix Option 160
Dynamic Addressing Methods with SLAAC and DHCPv6 162
Summary 162
Review Questions 163
References 164
Endnote 164
RFCs 164
Websites 165

Chapter 6 Link-Local Unicast Address 167

Structure of a Link-Local Unicast Address 169
Automatic Configuration of a Link-Local Address 170
EUI-64 Generated Interface ID 170
Verifying the Router's Link-Local Address on Ethernet and Serial Interfaces 174
Randomly Generated Interface ID 175
Zone ID (%) on Link-Local Interfaces 176
Manual Configuration of a Link-Local Address 179
Link-Local Address and Duplicate Address Detection 182
Link-Local Addresses and Default Gateways 183
ipv6 enable: Isolated Link-Local Address 184
Pinging a Link-Local Address 186
Summary 189
Review Questions 190
References 191
RFCs 191

Chapter 7 Multicast Addresses 193

Scope 195
Multicast with Link-Local Scope Versus Link-Local Unicast Addresses 197
Well-Known Multicast Addresses 198
Solicited-Node Multicast Addresses 202
Mapping Unicast Address to Solicited-Node Multicast Address 204
Mapping to the Ethernet MAC Address 206
Mapping Solicited-Node Multicast to Ethernet MAC Addresses 206
Mapping Well-Known Multicast to Ethernet MAC Addresses 210
Verifying the Address Mappings on Cisco IOS, Windows, and Linux 210

Multiple Devices Using the Same Solicited-Node Multicast Address 212
One Solicited-Node Multicast Address for Multiple Unicast Addresses 214
Multicast Listener Discovery 216
MLD Snooping 220
Summary 221
Review Questions 222
References 222
RFCs 222
Websites, Videos, and Books 223

Part III Dynamic IPv6 Addressing 225

Chapter 8 Basics of Dynamic Addressing in IPv6 227
Dynamic IPv4 Address Allocation: DHCPv4 227
Dynamic IPv6 Address Allocation 229
ICMPv6 Router Solicitation and Router Advertisement Messages 230
Router Advertisement Methods and the A, O, and M Flags 233
Method 1: Stateless Address Autoconfiguration (SLAAC) 235
Method 2: SLAAC with Stateless DHCPv6 237
Method 3: Stateful DHCPv6 238
DHCPv6 Services 240
DHCPv6 Terminology and Message Types 241
DHCPv6 Communications 245
Summary 248
Review Questions 249
References 250
RFCs 250
Website 250

Chapter 9 Stateless Address Autoconfiguration (SLAAC) 251
The RA Message and SLAAC 252
On-Link Determination 258
Generating an Interface ID 260
Generating the Interface ID Using the EUI-64 Process 261
Configuring a Windows Host to Use EUI-64 264
Privacy Extension for Stateless Address Autoconfiguration 266
Privacy Extension and Generating Randomized Interface IDs 267
Privacy Extension and Temporary Addresses 268

Disabling the Use of Temporary Addresses 269
Autoconfigured Address States and Lifetimes 270
Example: Autoconfigured Address States and Lifetimes 272
Displaying IPv6 Lifetimes and State Information on Windows, Linux, and Mac OS 278
Router Advertisement Fields and Options 279
Examining the Router Advertisement with Wireshark 279
Modifying the Valid Lifetime and Preferred Lifetime in the RA Message 282
Including the DNS Address in the Router Advertisement 282
Router Advertisement Configuration Options 284
Default Address Selection 288
Configuring the Router's Interface as a SLAAC Client 290
Summary 290
Review Questions 292
References 294
RFCs 294
Websites 295
Other 295

Chapter 10 Stateless DHCPv6 297

SLAAC with Stateless DHCPv6 298
Implementing Stateless DHCPv6 300
Configuring the RA Message's Other Configuration Flag 300
Wireshark Analysis of Router Advertisement: SLAAC and Stateless DHCPv6 301
Configuring a Router as a Stateless DHCPv6 Server 303
Verifying Stateless DHCPv6 on a Windows Client 304
Verifying the Router as a Stateless DHCPv6 Server 305
DHCPv6 Options 306
rapid-commit Option 306
Configuring the Rapid-Commit Option 307
Relay Agent Communications 308
DHCPv6 Relay Agent Configuration Commands 310
Configuring a Unicast DHCPv6 Relay Agent 311
Configuring a DHCPv6 Relay Agent Using a Multicast Address 311
Summary 312
Review Questions 313

References 314
RFCs 314
Websites 314

Chapter 11 Stateful DHCPv6 315

Stateful DHCPv6 Messages and Process 316
Implementing Stateful DHCPv6 317
Configuring the RA Message M Flag and A Flag 318
Setting the M Flag to 1 with an A Flag Set to 1 318
Consequences of Disabling the RA Message or Omitting the Prefix 319
Setting the M Flag to 1 and Modifying the A Flag to 0 320
Wireshark Analysis of Router Advertisement: Stateful DHCPv6 322
Configuring a Router as a Stateful DHCPv6 Server 323
The ***Address Prefix*** *Command 325*
Verifying Stateful DHCPv6 on a Windows Client 326
Verifying the Router as a Stateful DHCPv6 Server 327
DHCPv6 Options 329
IPv6 Prefix Delegation Options for DHCPv6 329
Sample Configuration: Prefix Delegation with DHCPv6 331
DHCPv6-PD Process 331
HOME Router (Requesting Router) Configuration and Verification 333
ISP Router (Delegating Router) Configuration and Verification 337
Verifying Prefix Delegation with DHCPv6 on WinPC 339
Summary 340
Review Questions 341
References 343
RFCs 343
Websites 343

Part IV ICMPv6 and ICMPv6 Neighbor Discovery 345

Chapter 12 ICMPv6 347

General Message Format 348
ICMP Error Messages 352
Destination Unreachable 352
Packet Too Big 355
Path MTU Discovery 355

Time Exceeded 357
Parameter Problem 360
ICMP Informational Messages 361
Echo Request and Echo Reply 361
Pinging a Global Unicast Address 362
Pinging a Link-Local Address 365
Summary 368
Review Questions 369
References 371
RFCs 371

Chapter 13 ICMPv6 Neighbor Discovery 373
Neighbor Discovery Options 374
Default Router and Prefix Determination 375
Router Solicitation Message 375
Router Advertisement Message 378
Address Resolution 384
The Address Resolution Process 385
Characteristics of the Neighbor Solicitation Message 388
Format of the Neighbor Solicitation Message 391
Format of the Neighbor Advertisement Message 393
Neighbor Cache 396
Destination Cache 401
Duplicate Address Detection (DAD) 402
Neighbor Unreachability Detection (NUD) 404
Redirect Message 405
Summary 407
Review Questions 408
References 411
RFCs 411

Part V Routing IPv6 413

Chapter 14 IPv6 Routing Table and Static Routes 415
Configuring a Router as an IPv6 Router 416
Understanding the IPv6 Routing Table 418
Codes: NDp and ND 420
Code: Connected 422
Code: Local 423

Configuring IPv6 Static Routes 424
Static Routes with a GUA Next-Hop Address 426
Static Routes with a Link-Local Next-Hop Address 427
Static Routes with Only an Exit Interface 428
Default Static Routes with Link-Local Next-Hop Addresses 429
Verifying IPv6 Static Routes 430
Summarizing IPv6 Routes 433
IPv6 Summary Static Route 435
CEF for IPv6 436
Summary 438
Review Questions 439
References 441
RFCs 441
Websites 441
Books 441

Chapter 15 EIGRP for IPv6 443
Comparing EIGRPv4 and EIGRPv6 444
Classic EIGRP for IPv6 446
Configuring Classic EIGRP for IPv6 447
Verifying Classic EIGRP for IPv6 450
EIGRP Named Mode for IPv6 456
Configuring EIGRP Named Mode for IPv6 457
Verifying EIGRP Named Mode for IPv6 464
Comparing EIGRP Named Mode for IPv4 and IPv6 468
Summary 470
Review Questions 472
References 473
RFC 473
Websites 473
Books 473

Chapter 16 OSPFv3 475
Comparing OSPFv2 and OSPFv3 476
Traditional OSPFv3 479
Configuring Traditional OSPFv3 480
ASBR and Advertising a Default Route 481
Area Border Router with Totally Stubby Area 482

Internal Router: Totally Stubby Area 483
Advertising a Default Route 484
Verifying Traditional OSPFv3 485
OSPFv3 with Address Families 492
Configuring OSPFv3 with AF 493
ASBR and Advertising a Default Route 493
ABR with Totally Stubby Area 497
Internal Router: Totally Stubby Area 498
Verifying OSPFv3 with AF 499
Configuring OSPFv3 for an IPv4 Island 507
Summary 509
Review Questions 511
References 513
RFCs 513
Websites 513
Books 513

Part VI Implementing IPv6 515

Chapter 17 Deploying IPv6 in the Network 517
IPv6 Address Plan Considerations 518
Encoding Information in the Subnet ID 521
VLAN-Mapped Subnet ID 523
IPv6 Address Plans 524
IPv6 VLANs 525
IPv6 First Hop Redundancy Protocols 529
ICMPv6 Neighbor Discovery 530
HSRP and VRRP 533
GLBP 534
Selecting an FHRP 536
Dual Stack 536
IPv6 Address Format in URL Syntax 538
DNS 539
DNS Query and Response 543
Happy Eyeballs 545
IPv6 Access Control Lists 546
Configuring IPv6 ACLs 546
Transition Technologies 550

Translation with NAT64 551
Traffic Initiated from IPv6-Only Clients to IPv4-Only Servers 553
Traffic Initiated from IPv4-Only Clients to IPv6-Only Servers 557
Other Translation Techniques 559
Tunneling IPv6 560
Conclusion 566
Summary 566
Review Questions 568
References 570
RFCs 570
Websites 571

Appendixes

Appendix A Configuring NAT64 and IPv6 Tunnels 573
Configuring NAT64 573
Configuring IPv6 Tunnels 577
Manual Tunnels 577
6to4 Tunnels 584
6to4 Tunnels and Loopback Interfaces 592
ISATAP 593

Appendix B IPv6 Command Quick Reference 601
Cisco IOS Commands 601
Addressing Commands 601
Global Unicast Address and Unique Local Unicast Addresses 601
Link-Local Unicast Address 601
General Prefix 602
DNS host commands 602
Verifying Address Information 602
ICMPv6 Router Advertisement Commands 602
Enabling ICMPv6 Router Advertisements 602
Modifying Router Advertisement Parameters on the Interface 602
Verifying Router Advertisements 603
Configuring a DHCPv6 Server 604
Stateless DHCPv6 Configuration Pool Commands 604
Stateful DHCPv6 Configuration Pool Commands 604
Associating the DHCPv6 Pool to an Interface 604

DHCPv6 Relay 605
Verifying DHCPv6 Information 605
IPv6 Access Control Lists 605
Configuring IPv6 ACLs 605
Verifying IPv6 ACLs 605
Static Routes, Displaying the Routing Table, and CEF for IPv6 605
Static Routes 605
Verifying Static Routes 606
CEF for IPv6 606
EIGRP for IPv6 606
Classic EIGRP for IPv6 606
EIGRP Named Mode 607
EIGRP for IPv6 Verification Commands 607
OSPFv3 608
Configuring Traditional OSPFv3 608
Verifying Traditional OSPFv3 609
Configuring OSPFv3 with Address Families 609
Verifying OSPFv3 with Address Families 610
Host Operating System Commands 610
Windows OS 610
General Commands 610
Interface Addresses Information 611
SLAAC Interface ID 611
Linux OS 612
General Commands 612
Address Configuration Commands 613
Mac OS X 613
General Commands 613
Address Configuration Commands 614

Appendix C Answers to Review Questions 615

Index 631

Icons Used in This Book

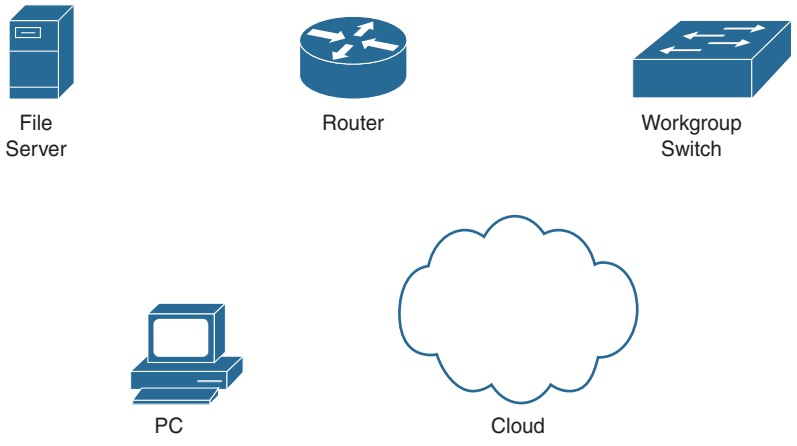

Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italics* indicate arguments for which you supply actual values.
- Vertical bars | separate alternative, mutually exclusive elements.
- Square brackets [ ] indicate optional elements.
- Braces { } indicate a required choice.
- Braces within brackets [{ }] indicate a required choice within an optional element.

Introduction

This book is intended to give you an in-depth understanding of IPv6 and related protocols. This book is for those who are new to the networking field, such as computer networking students, and also for network engineers with many years of experience administering larger enterprise networks. The only prerequisite is a basic understanding of networking protocols, including IPv4.

This second edition is a complete reorganization and almost a complete rewrite of the first edition, with a lot of new content. I have been a teacher for over 20 years, and this book is written for students of IPv6, as a self-study guide for learning IPv6. This book is designed to walk you through learning about IPv6 as if you were in a classroom, with the instructor explaining each concept. The information has been organized to help both those who want to read the book from cover to cover and also for those looking for specific information.

There is a great deal to learn about IPv6, and it is much more than just becoming familiar with a larger address. A brief look at the contents of this book will give you an idea of what is covered and what is needed to have a good understanding of IPv6.

My approach to writing this book was to do my best to explain each concept in a simple, step-by-step approach, as well as to include the critical details. It was a challenging balance between providing as much information as possible and not overwhelming you, the reader. IPv6 is not difficult to learn but involves multiple protocols and processes that might be new to some.

Don't be overwhelmed by all the details. For example, although I have included a brief description of each field in the protocols discussed in the book, it isn't necessarily important that you understand the details of each one. I mention this throughout the book. But I did feel it necessary not to leave out or hide these details from you.

RFCs are cited throughout the book. It was important to include these references for two reasons. First of all, I wanted to give you the authoritative source for the material in this book so that you have resources for more information. Second, IPv6 is currently and will continue to be a moving target for quite some time. Although it has been around for many years, additional development and fine-tuning are still taking place. If you are not familiar with reading RFCs, don't be intimidated. Most of them are not difficult to read, and they do their best to explain the topic clearly.

Review questions are included at the end of each chapter to help you understand some of the fundamental concepts discussed in the chapter. The review questions provide a high-level overview of some the key points discussed in the chapter. They are not meant as a detailed assessment of all the material covered in the chapter. An "IPv6 Command Quick Reference" has been included as Appendix B for Cisco IOS, Windows, Linux, and Mac OS X commands.

At times in this book I introduce a technology or concept but state that it is covered in more detail in a later chapter. I do this to explain the concept as it relates to the topic being discussed without getting lost in the details. The details are covered where appropriate. I feel it is better to revisit some of the more advanced topics after you have

a more complete understanding of the entire IPv6 topic. At times I state that a concept is "beyond the scope of this book." I suggest resources for those who might be interested in learning more about those topics.

The objective of this book is to explain IPv6 as clearly as possible. At times it was like herding cats, trying to decide which topic to cover first. Chapter 2 is an IPv6 primer designed to give an overview of the main topics. Having this overview will make it easier as you progress through the rest of the book.

Readers are welcome to use the resources on my website, www.cabrillo.edu/~rgraziani, for IPv6, CCNA, or CCNP information. You can email me at graziani@cabrillo.edu to obtain the username and password for all my materials.

Goals and Methods

The most important goal of this book is to provide a thorough yet easy-to-understand introduction to IPv6. It is written for both computer networking students and seasoned network engineers. This book is also intended to provide a foundation in IPv6 that will allow you to build on. It explains topics that might be a little more challenging to grasp.

Another goal of this book is to be a resource for IPv6. The book is organized to make it as easy as possible to find information on specific topics. I have included command syntax, RFCs, and links to Cisco white papers to help guide you toward a better understanding of many of the topics.

Who Should Read This Book

This book is intended for anyone seeking a solid understanding of the fundamentals of IPv6, such as network engineers, network designers, network technicians, technical staff, and networking students, including those who are part of Cisco Networking Academy. You should have a basic familiarity with IPv4 and networking protocols before you begin reading this book.

Professionals deploying or planning to deploy IPv6 within a network will find this book useful. It provides examples, figures, IOS commands, and recommendations for configuring Cisco IOS IPv6 technology. Although Cisco devices are used in this book, those using non-Cisco equipment will also find this book helpful. The vast majority of protocols and technologies are IETF standards. Configuration and verification commands for Windows, Linux, and Mac OS are also included throughout the book.

How This Book Is Organized

If you are new to IPv6, you should read this book from cover to cover. However, if you have some knowledge of IPv6, it is designed to be flexible and allows you to easily move between chapters and sections of chapters to cover just the material you want to review. A common topology is used throughout the book except in a few cases.

Chapters 1 through 3 provide an introduction to IPv6, the reasons for moving to IPv6, an IPv6 primer, and a comparison of the IPv4 and IPv6 protocols. Chapters 4 through 7 discuss the different types of IPv6 address, including how to represent IPv6 addresses, global unicast addresses, link-local unicast addresses, and IPv6 multicast addresses. Chapters 8 through 11 discuss dynamic IPv6 addressing methods. Dynamic address allocation differs significantly in IPv6 and IPv4. These chapters discuss Stateless Address Autoconfiguration (SLAAC), stateless DHCPv6, and stateful DHCPv6. The chapter on SLAAC discusses the reason for permanent and temporary addresses and how to manage them using Cisco IOS and host operating systems. These chapters include the Cisco IOS commands and configuration examples for stateless and stateful DHCPv6. Chapters 12 and 13 discuss ICMPv6 and ICMPv6 Neighbor Discovery Protocol. These protocols and messages are introduced in earlier chapters, beginning with Chapter 2. Chapters 12 and 13 examine ICMPv6 and ICMPv6 Neighbor Discovery in more detail. Chapters 14 through 16 cover routing IPv6, including the IPv6 routing table, classic EIGRP for IPv6, EIGRP named mode, traditional OSPFv3, and OSPFv3 with address families. The last chapter, Chapter 17, introduces deploying IPv6 and transitioning from IPv4 to IPv6. In case you intend to read all the chapters, the order in the book is sequential.

The following list highlights the topics covered in each chapter and the book's organization:

- **Chapter 1, "Introduction to IPv6":** This chapter discusses how the Internet of today requires a new network layer protocol, IPv6, to meet the demands of its users. It also examines the limitations of IPv4 and describes how IPv6 resolves these issues while offering other advantages as well. This chapter examines the rationale of IPv6 and concerns regarding IPv4 address depletion. It presents a brief history of both IPv4 and IPv6. The IPv4 migration technologies CIDR and NAT are also discussed.
- **Chapter 2, "IPv6 Primer":** This chapter introduces some of the basic concepts and protocols that are explained in more detail throughout the rest of the book, including IPv6 address types, the basics of dynamic address allocation, and the hexadecimal number system, which is used to represent IPv6 addresses. This chapter gives an overview of some IPv6 concepts that are helpful in learning IPv6. This chapter also highlights many of the differences in IPv6.
- **Chapter 3, "Comparing IPv4 and IPv6":** This chapter compares and contrasts the IPv4 and IPv6 protocols. It also explores how fragmentation is handled. It discusses the IPv6 extension headers as well.
- **Chapter 4, "IPv6 Address Representation and Address Types":** This chapter introduces IPv6 addressing and address types. It discusses representation of IPv6 addresses, along with the different formats for representing IPv6 addresses and the rules for compressing the IPv6 notation. This chapter provides an introductory look at the different types of IPv6 addresses, including unicast, multicast, and anycast. It also discusses prefix length notation.
- **Chapter 5, "Global Unicast Address":** This chapter examines the global unicast address in detail. It examines the different parts of a global unicast address as well as

manual configuration of a global unicast address for Cisco IOS and host operating systems. The chapter also covers subnetting of IPv6, along with prefix allocation.

- **Chapter 6, "Link-Local Unicast Address":** This chapter examines link-local addresses and includes static and dynamic link-local address configuration examples. It explains the EUI-64 process, along with the significance of a link-local address in IPv6.
- **Chapter 7, "Multicast Addresses":** This chapter examines multicast addresses, including well-known and solicited-node multicast. It discusses the advantages of a multicast address over a broadcast address (the broadcast address does not exist in IPv6) and how IPv6 multicast addresses are mapped to Ethernet MAC addresses.
- **Chapter 8, "Basics of Dynamic Addressing in IPv6":** This chapter introduces and compares the three methods of dynamic address allocation: Stateless Address Autoconfiguration (SLAAC), stateless DHCPv6, and stateful DHCPv6. The chapters that follow discuss these methods in more detail.
- **Chapter 9, "Stateless Address Autoconfiguration (SLAAC)":** This chapter discusses the SLAAC process in detail. It includes a Wireshark examination of an ICMPv6 Router Advertisement message suggesting SLAAC. This chapter discusses the use of the privacy extension and temporary addresses with SLAAC-generated addresses, including the different states and lifetimes. It also explains how to manage privacy options on host operating systems.
- **Chapter 10, "Stateless DHCPv6":** This chapter examines SLAAC and other stateless DHCPv6 services. It covers DHCPv6 terminology and message types, along with the DHCPv6 process between the client and server. This chapter explains the rapid-commit option and relay agents.
- **Chapter 11, "Stateful DHCPv6":** This chapter examines stateful DHCPv6 services, similar to DHCP for IPv4. It also introduces a common method for providing IPv6 address space to homes using DHCPv6 with Prefix Delegation.
- **Chapter 12, "ICMPv6":** This chapter examines ICMPv6, which is a much more robust protocol than ICMPv4. It covers ICMPv6 error messages, including Destination Unreachable, Packet Too Big, Time Exceeded, and Parameter Problem. It also covers the ICMPv6 Echo Request and Echo Reply informational messages, along with Multicast Listener Discovery messages.
- **Chapter 13, "ICMPv6 Neighbor Discovery":** This chapter examines ICMPv6 Neighbor Discovery, including Router Solicitation, Router Advertisement, Neighbor Solicitation, Neighbor Advertisement, and Redirect messages. Not only does IPv6 resolve larger address space issues but ICMPv6 with Neighbor Discovery Protocol also presents a major change in network operations, including link-layer address resolution (ARP in IPv4), Duplicate Address Detection (DAD), Stateless Address Autoconfiguration (SLAAC), and Neighbor Unreachability Detection (NUD). This chapter discusses the IPv6 neighbor cache and neighbor cache states, similar to those of the IPv4 ARP cache.

- **Chapter 14, "IPv6 Routing Table and Static Routes":** This chapter examines the IPv6 routing table. It also discusses the configuration of IPv6 static routes, which are similar to static routes for IPv4. It explains IPv6 default routes and route summarization, as well as CEF for IPv6.
- **Chapter 15, "EIGRP for IPv6":** This chapter discusses EIGRP for IPv6. It begins with a comparison of EIGRP for IPv4 and EIGRP for IPv6. It discusses configuration and verification of classic EIGRP for IPv6 and EIGRP named mode (for IPv4 and IPv6).
- **Chapter 16, "OSPFv3":** This chapter discusses OSPFv3. It begins with a comparison of OSPFv2 (IPv4 only), traditional OSPFv3 (IPv6 only), and OSPFv3 with address families (IPv4 and IPv6). It also discusses configuration and verification of traditional OSPFv3 and OSPFv3 with address families.
- **Chapter 17, "Deploying IPv6 in the Network":** This chapter covers strategies for deploying IPv6, including creating an IPv6 address plan, configuring IPv6 VLANs, and implementing transparent failover at the first-hop router, using ICMPv6 or a first-hop redundancy protocol (FHRP). This chapter also discusses IPv4 and IPv6 integration and coexistence, including dual stacking, NAT64, and tunneling.
- **Appendix A, "Configuring NAT64 and IPv6 Tunnels":** This appendix provides configuration examples and additional information on NAT64 and IPv6 tunnels, introduced in Chapter 17.
- **Appendix B, "IPv6 Commands Quick Reference":** This appendix provides a summary of the Cisco IOS, Windows, Linux, and Mac OS commands used in this book.
- **Appendix C, "Answers to Review Questions":** This appendix provides the answers to the Review Questions at the end of each chapter.

This page intentionally left blank

Part I

Introduction to IPv6

Chapter 1 Introduction to IPv6

Chapter 2 IPv6 Primer

Chapter 3 Comparing IPv4 and IPv6

This page intentionally left blank

Chapter 1

Introduction to IPv6

"I am a little embarrassed about that because I was the guy who decided that 32-bit was enough for the Internet experiment. My only defense is that that choice was made in 1977, and I thought it was an experiment. The problem is the experiment didn't end, so here we are."

—Vint Cerf, speaking in 2011, credited as one of the co-founders of the Internet

Internet Protocol version 6 (IPv6), the successor to Internet Protocol version 4 (IPv4), has been a long time coming. For more than 35 years, IPv4 has been an integral part of the Internet evolution. But it was designed at a time when there were fewer than 600 hosts on the Internet, before the World Wide Web, email, video streaming, mobile phones, and countless other innovations. This chapter discusses the limitations of IPv4 and how after many years of proclaiming the need for IPv6, IPv6 is really finally here.

IPv6 Is Here

IPv4 provides for a maximum 4.29 billion 32-bit addresses, which seemed like more than enough addresses when IPv4 became a standard in 1980. There were about 4.5 billion people on Earth at the time, so even if every person on the planet needed one IPv4 address, it seems like that should have been enough.

The number of IP addresses needed today far exceeds the world's population for several reasons. First of all, IPv4 addresses are often allocated in groups of addresses, as network addresses. Places such as companies, schools, homes, cafés, and airports are allocated network addresses for their users. As I go from my home to the café, or to my college, I receive a different IP address to access each network. So, in most cases, I need a different IPv4 address wherever I go.

But a much larger reason we need so many more IP addresses is the number of devices per person that are being connected to the Internet—and the number of people who still need access to the Internet. A lot of the devices that have Internet connections are obvious: computers, laptops, mobile phones, and tablets. But then there are also devices that many people don't realize have IP addresses, such as cable or satellite subscriber set-top boxes and cars. Beginning in 2016, statistics compiled by Chetan Sharma Consulting showed that U.S. Internet service providers (ISPs) connected more new cars to their networks than they did new mobile phones.[1] We will examine this more closely later in this chapter.

IPv6 provides more addresses than IPv4. As mentioned earlier, IPv4, with its 32-bit address space, provides for 4,294,967,296, or 2^{32}, addresses. In comparison, the 128-bit IPv6 address space provides for 340 undecillion addresses! How large is 340 undecillion? It equates to 340,282,366,920,938,463,463,374,607,431,768,211,456, or 2^{128}. As you can see, IPv6 provides a tremendous—almost unimaginable—number of addresses.

But IPv6 is much more than just a lot of addresses. The designers of IPv6 took the opportunity to improve IP and related protocols. Throughout this book we discuss the many changes and improvements made to IPv6 along with other protocols, such as ICMPv6.

Many people have been skeptical about the need for IPv6. The proponents of IPv6 have been preaching the need to move to IPv6 since the mid-1990s. And yet, after all these years, many people still don't see the need for IPv6.

However, after countless proclamations that the IPv4 sky is falling, a more mature IPv6 is finally being deployed at a growing rate. IPv6 is now enabled by default on every major host operating system, including Windows, Mac OS, and Linux. All mobile operating systems are also IPv6-enabled, including Google Android, Apple iOS, and Windows Mobile. Because of the enormous number of mobile devices, the mobile carriers are leading the way in IPv6 deployments. For example, World IPv6 Launch measurements showed that over 75% of Verizon Wireless's overall traffic is over IPv6[2]. T-Mobile phones are IPv6-only; they use a special protocol translation technique when communicating with IPv4-only devices.[3] Gaming stations such as Microsoft's Xbox One and Sony's PlayStation are also IPv6-enabled. IPv6 is here.

Note For information and presentation material from service providers, search for "IPv6" at www.nanog.org/archives/presentations, the North American Network Operators' Group (NANOG) meeting presentation archives.

Comcast, the world's largest media company, has enabled IPv6 throughout its network, which includes the majority of residences in the United States. Most Comcast residential customers don't realize that they may be using IPv6 when communicating with Google, Facebook, Netflix, LinkedIn, YouTube, and many other content providers. As of this writing, more than 50% of Internet web content is available over IPv6. The four major

US mobile providers—Verizon Wireless, T-Mobile USA, Sprint Wireless, and AT&T Wireless—send the majority of their traffic over IPv6 to major IPv6-capable content providers.

IPv6 is here and is being adopted…finally. Given all these pronouncements over the years, it's easy to ignore seeing the need to transition to IPv6. Since 1996, we have been hearing how we must begin to migrate to IPv6, yet the IPv4 Internet has continued to survive and even grow.

What's changed now, so that we have to transition to IPv6? It's simple: We're out of IPv4 addresses. The need for IP addresses is real. It is predicted that the Internet of Things (IoT)—everyday objects having Internet connectivity—will add about another 35 billion devices by 2020. Cloud services such as Amazon AWS and Microsoft Azure now support IPv6. About half of the world's population is yet to get Internet access, particularly in Africa and Asia. Already about 5.5 billion mobile phones in the world require Internet connectivity; and remember that IPv4 allows for only 4.29 billion addresses. We simply don't have the IPv4 address space to accommodate current needs—let alone future growth. (We'll talk more about the details of the depletion of IPv4 address space and the need for more addresses later in this chapter.)

For those who haven't already started the transition to IPv6, now is a good time to begin to become familiar with IPv6. If you're a network administrator, your Windows, Mac, Linux, Android, and iOS clients are already running IPv6. At the very least, it is necessary to understand and secure your IPv6 network to prevent man-in-the-middle (MITM) and denial-of-service (DoS) attacks.

The following sections discuss issues including the following:

- Why transition to IPv6?
- "Short-term" solutions to IPv4 address exhaustion: CIDR and NAT with private IPv4 addresses
- What happened to IPv5?
- The development history of IPv6
- Some of the myths about IPv6
- Transitioning to IPv6

Why Transition to IPv6?

What are the advantages of using IPv6? Really, who cares? If you're a network administrator, you might be thinking, "I have enough work to do, and my IPv4 network works just fine." And for now, you may still have plenty of IPv4 addresses available. However, continuing to ignore IPv6 will begin to have a real business impact on your network. At the very least, ignoring IPv6 could put your network at risk.

The following sections discuss these benefits of IPv6 and reasons to begin the transition to IPv6:

- IPv4 address depletion
- Access to IPv6-only customers
- Better performance
- Securing your current network

IPv4 Address Depletion

There is no getting around it. We are running out of IPv4 address space. Sooner or later, network operators will need to transition their networks to IPv6. Even if you're not ready to implement IPv6 on your network right now, there are many reasons you need to understand IPv6.

As mentioned earlier, choosing to ignore IPv6 as IPv4 exhaustion continues may put your business and your network at risk. IPv4 address depletion is discussed in more detail later in this chapter.

Access to IPv6-Only Customers

Already, there are places in the world that are IPv6-only. In some areas of the world today, it is simply not possible to get an IPv4 address. This will become more and more the norm as IPv4 address space becomes even less available. Unless you want to risk isolating your business from these users, it is going to become increasingly important for your network to be enabled for IPv6.

Yes, there are translation techniques that allow IPv4-only and IPv6-only networks to communicate with each other, but they are not always reliable, and they are often accompanied by degraded performance. Also, any sort of translation mechanism makes it difficult to measure the quality of the user experience and has the potential to break certain applications.

Better Performance

Increased performance can be another benefit of transitioning to IPv6. Many content providers are seeing substantial increases in performance with IPv6. Facebook sees 20% to 40% better performance of news feeds with IPv6. Tests at Time Warner Cable show about a 15% performance increase with IPv6. Asia-Pacific Network Information Centre (APNIC) says that IPv6 traffic is faster over trans-Pacific links due to better routing aggregation.[4]

Studies on why IPv6 may be faster than IPv4 are still incomplete. However, the Internet Society (ISOC) has stated, "The reality is that connections from users on IPv6 networks to services that run over IPv6 *are going to be faster* than connections to services still on legacy IPv4 networks that have to go through middleboxes such as NAT devices or application-layer gateways (ALGs)."[5]

As discussed in more detail later in this chapter, network address translation (NAT) was never intended to be a permanent solution to the IPv4 exhaustion problem. It was meant as a temporary solution. Unfortunately, NAT breaks many applications and adds latency.

Securing Your Current Network

As mentioned previously, many home subscribers are seeing performance benefits thanks to IPv6. This is because residential service providers such as Comcast have enabled IPv6 to all their residential customers. What have customers had to do to enable their computers and mobile devices for IPv6? Nothing. This is because IPv6 has been enabled on all major operating systems—Windows, Mac OS, Linux, iOS, and Android—for years.

Even if your provider isn't running IPv6, your host operating systems are. Example 1-1 shows a Windows host running both IPv4 and IPv6. This is known as *dual stack*, meaning a device is running both the IPv4 and IPv6 protocols at the same time. Notice in Example 1-1 that the Windows host has a link-local IPv6 address. (Link-local addresses are explained Chapter 4, "Link-Local Unicast Address.") An IPv6 device automatically creates an IPv6 link-local address during startup. So without requiring the user to do anything, this Windows device is enabled and running IPv6.

Example 1-1 *Windows Running IPv6 by Default*

```
PC> ipconfig
Windows IP Configuration
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  : localdomain
   Link-local IPv6 Address . . . . : fe80::50a5:8a35:a5bb:66e1%7
   IPv4 Address. . . . . . . . . . : 192.168.1.101
   Subnet Mask . . . . . . . . . . : 255.255.255.0
   Default Gateway   . . . . . . . : 192.168.1.1
```

Note If the **ipconfig** command displays **IPv6 Address** without the preceding **Link-local**, it is a *global unicast address*, which is an Internet-routable IPv6 address. The **ifconfig** command shows similar output on Mac OS and Linux devices.

Because host operating systems are enabled for IPv6 by default, it is critical for network administrators to become familiar with IPv6. A network may be secured against DoS, MITM, and other attacks for IPv4, and it's important to provide the same security measures for IPv6. Host operating systems are ready and willing to use IPv6 to communicate, even if you're not.

IPv4

Before discussing IPv6 in more detail, let's look at the history of IPv4. IPv4 is the current Layer 3 protocol used on the Internet and most other networks. IPv4 has been an integral part of the Internet evolution over the past 35 years. It was originally described in RFC 760 (January 1980) and replaced by RFC 791 (September 1981).

Can you imagine what the Internet was like in those early years? IPv4 was introduced when there were fewer than 600 hosts on the Internet, so its address space of 4.29 billion seemed enormous. Although many of the same concepts of packet switching apply today, obviously the number of users on the Internet and how the Internet is used are vastly different than they were back then. Even with the advent of the World Wide Web in the early 1990s, there were only about 16 million users on the Internet worldwide; in 2016 there were more than 3.6 billion users. And remember that in the early 1990s, there was about one device per user. Today each user has multiple devices. The total number of devices connected to the global Internet today is 15 billion.

The reasons to migrate from IPv4 to IPv6 are discussed later in this chapter. For now, just recognize that presently the Internet is a much different place than it was 30, 10, or even 5 years ago. Today's Internet is much more than just web pages, email, and file transfers. With the explosive growth of mobile devices and peer-to-peer networking, along with the potential impact of Internet-ready consumer goods, we have gone from just an Internet of computers to an *Internet of Things* (*IoT*). Originally the idea of the IoT focused on objects having radio-frequency identification (RFID) tags for identification and inventory purposes. It has since evolved into something much larger, including smart grids, smart and connected cities, home automation, automotive computing, and more. A new and growing trend in the IoT is wearable devices that do everything from measuring your heart rate to projecting an image of your smart phone on your arm.

IPv4 Address Depletion

> *"IPv4 depletion is THE reason to make sure you have a plan to deploy IPv6… IPv6 implementation is the only viable solution to the dwindling pool of IPv4 space."*
>
> —American Registry for Internet Numbers (ARIN)

The number of people accessing the Internet is increasing dramatically. Even with short-term solutions like NAT, we are in the final stages of public IPv4 address availability. Table 1-1 provides some world Internet usage and population statistics, as of June 30, 2016, that reveal some interesting facts regarding IPv4 address allocation. Notice that North America has a penetration rate of about 89%. This means that the vast majority of people in North America use the Internet. Asia, on the other hand, has a penetration rate of only about 46%. The problem becomes apparent when you compare the two populations. Whereas North America has a population of about 360 million, Asia has 4 billion people. As Asia, Africa, and other areas of the world become increasingly connected to the Internet, there is no way IPv4 can keep up, even with NAT.

Table 1-1 *World Internet Usage and Population Statistics as of June 30, 2016*

World Region	Population (2016 Est.)	Population % of World	Number of Internet Users	Penetration Rate	Growth 2000–2016
Africa	1,185,529,578	16.2%	340,783,342	28.7%	7,448.8%
Asia	4,052,652,889	55.2%	1,846,212,654	45.6%	1,515.2%
Europe	832,073,224	11.3%	614,979,903	73.9%	485.2%
Latin America/ Caribbean	626,119,788	8.5%	384,751,302	61.5%	2,029.4%
Middle East	246,700,900	3.4%	141,489,765	57.4%	4,207.4%
North America	359,492,293	4.9%	320,067,193	89.0%	196.1%
Oceania/ Australia	37,590,820	0.5%	27,540,654	73.3%	261.4%
World total	7,340,159,492	100.0%	3,675,824,813	50.1%	918.3%

Note: The Internet penetration rate is the percentage of the population of a country or region that uses the Internet.

Source: The information in this table was obtained from www.internetworldstats.com, copyright © 2001–2016, Miniwatts Marketing Group. All rights reserved worldwide.

There has also been extraordinary growth in the number of devices connected to the Internet. GlobalWebIndex estimates that consumers own 3.64 devices per person, and this number continues to increase every year.[6] When you start tallying up the Internet media devices, smart TVs, smart home devices, wearables, and other IoT devices, you can see the phenomenal number of devices needing Internet connections. This number is, again, much larger than IPv4 can accommodate.

The Internet Assigned Numbers Authority (IANA) assigns IPv4 addresses to the five Regional Internet Registries (RIRs) in /8 address blocks. (IANA is also responsible for allocating IPv6 address space to RIRs.) Using this address space, RIRs then redistribute the IPv4 network addresses to ISPs and other end customers. The five RIRs and the areas they manage are as follows:

- **African Network Information Centre (AFRINIC):** Africa
- **American Registry for Internet Numbers (ARIN):** United States, Canada, some parts of the Caribbean region, and Antarctica
- **Asia-Pacific Network Information Centre (APNIC):** Asia, Australia, New Zealand, and neighboring countries
- **Latin America and Caribbean Network Information Centre (LACNIC):** Central America, South America, and most of the Caribbean region
- **Réseaux IP Européens (RIPE) Network Coordination Centre:** Europe, the Middle East, and Central Asia

On January 31, 2011, IANA allocated the last two blocks of IPv4 address space, 39.0.0.0/8 and 106.0.0.0/8, to APNIC, the RIR for the Asia-Pacific region. This triggered an IANA policy that the remaining five /8 addresses would be distributed equally among the five RIRs. At this point, IANA had now run out of IPv4 addresses.

On September 24, 2015, ARIN, the RIR for North America, had officially run out of IPv4 addresses. Currently, four out of the five RIRs have no more IPv4 addresses to allocate. AFRINIC is the last RIR to have any IPv4 addresses.[7] Although AFRINIC still has plenty of IPv4 address space, it is urging its subscribers to service both IPv6 and IPv4 to avoid lagging behind the rest of the world.[8]

Does this mean that IPv4 addresses are no longer available for end customers? No. Customers can still get IPv4 addresses from most ISPs. However, many ISPs are severely limited. In the wake of all this IPv4 address depletion, a "gray market" for IPv4 addresses has emerged. Several websites are serving as brokers for organizations that want to sell or lease IPv4 address space. ISOC has stated that buyers and sellers should make sure that any transfers use the appropriate RIR processes so that they don't compromise network operations or security.

Later in this chapter we will look at some of the benefits that IPv6 provides. However, the most compelling reason for moving to IPv6 is simply that we are out of IPv4 address. The "killer application" of IPv6 is the Internet itself—preserving the technologies of today and allowing for future growth and innovations.

There has been some inefficiency in the allocation of the 4.29 billion IPv4 addresses over the years. Even if it were possible to reassign the entire IPv4 address space in a more efficient manner, it would be a very short-term solution. IPv6 antagonists continue to find ways to extend the use of IPv4 without preparing to migrate to IPv6. It has often been said that this is a little like rearranging deck chairs on the *Titanic*. It might make things a little more comfortable for a while, but sooner or later, you must come to the realization that you're in trouble.

When Bob Kahn and Vint Cerf first developed the Transmission Control Protocol/Internet Protocol (TCP/IP) suite of protocols, including IPv4 and its 32-bit addresses, they never envisioned that the Internet would be what it is today—something that many rely on not just for communications but also for a variety of economic, social, political, and educational reasons.

The eventual exhaustion of IPv4 address space was recognized in the early 1990s, even before the World Wide Web came to be. And then the advent of the World Wide Web brought an explosion in the size and diversity of the Internet population. The people using the Internet were no longer limited to a relatively small group of computer-savvy users who could navigate this unknown global network using slightly cumbersome and less-than-intuitive tools such as File Transfer Protocol (FTP), newsgroups, and Unix-to-Unix Copy (UUCP). With Tim Berners-Lee's development of Hypertext Transfer Protocol (HTTP) at the European Organization for Nuclear Research (CERN) in Geneva, followed by the first web browser, Mosaic, created by Marc Andreessen and his team at the

National Center for Supercomputing Applications at the University of Illinois at Urbana–Champaign, all of a sudden, people who had never owned a personal computer were accessing the Internet. It was becoming clear that 4.29 billion addresses were not going to last long.

So, in the early 1990s, the Internet Engineering Task Force (IETF) began development of a new version of IP known as IP Next Generation, which later became IPv6. However, IPv6 was a long-term solution, and a short-term solution was needed immediately. Several short-term solutions were put into place. Two of the most important, discussed in the following sections, were Classless Inter-Domain Routing (CIDR) and NAT with private IPv4 addresses.

CIDR

In the early years, IPv4 network addresses were allocated using one of three classes—Class A, Class B, or Class C—as shown in Table 1-2.

Table 1-2 *Classful Unicast IPv4 Addresses*

Address Class	First Octet Range	Number of Possible Networks	Number of Hosts per Network
Class A	0 to 127	128 (2 are reserved)	16,777,214
Class B	128 to 191	16,384	65,534
Class C	192 to 223	2,097,152	254

There were a total of five different classes:

- **Class A:** Has a subnet mask of 255.0.0.0, or /8
- **Class B:** Has a subnet mask of 255.255.0.0, or /16
- **Class C:** Has a subnet mask of 255.255.255.0, or /24
- **Class D:** Used for multicast
- **Class E:** Experimental

Class A networks, with a /8 mask, can each have a total of 16,777,214 (2^{24}) hosts (with two addresses reserved for the network and broadcast addresses). The 128 (2^{7}) Class A networks and the enormous number of hosts per network account for 50% of the total IPv4 address space, as shown in Figure 1-1.

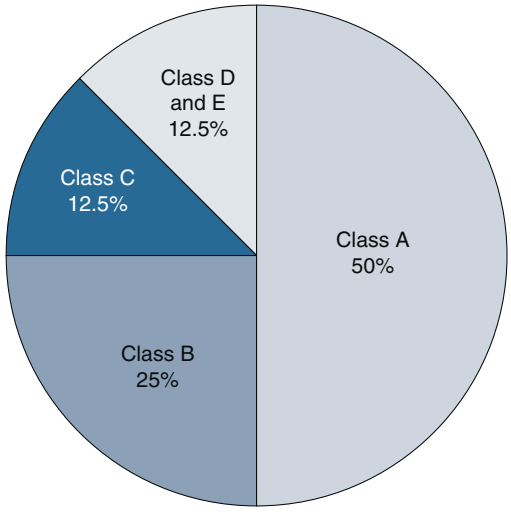

Figure 1-1 *Classful IPv4 Address Allocation*

Class B networks have a /16 mask, which means they can have a total of 65,534 (2^{16}) hosts per network. The 16,384 (2^{14}) Class B networks represent 25% of all IPv4 addresses.

Class C networks, with a /24 mask, can have a maximum of only 254 (2^{8}) hosts for each Class C network address. Even though there are 2,097,152 (2^{21}) Class C networks available, this represents only 12.5% of the total available IPv4 address space.

Because network administrators could obtain a network address from only one of these three classes, it made for a very inefficient method of allocating addresses. Most Class A and B networks had enormous numbers of unused addresses, while most Class C networks had comparatively very few.

In 1992, the IETF replaced this allocation method with Classless Inter-Domain Routing (CIDR, pronounced *cider*), which is described in RFC 1338 and later in RFC 1519. Network addresses were no longer allocated based on one of three classes. The addresses could be assigned with any size mask, as shown in Figure 1-2. (Note that Figure 1-2 shows the total number of possible host addresses and does not subtract two addresses for the network and broadcast addresses.) This is in contrast to the limitations of allocating a prefix length of just /8 (Class A), /16 (Class B), or /24 (Class C). The additional prefix lengths gave RIRs and ISPs more flexibility in allocating addresses to customers and made for a much more efficient distribution of the limited IPv4 address space.

Class A	**11111111.00000000.00000000.00000000**	**/8 (255.0.0.0)**	**16,777,216 host addresses**
	11111111.10000000.00000000.00000000	/9 (255.128.0.0)	8,388,608 host addresses
	11111111.11000000.00000000.00000000	/10 (255.192.0.0)	4,194,304 host addresses
	11111111.11100000.00000000.00000000	/11 (255.224.0.0)	2,097,152 host addresses
	11111111.11110000.00000000.00000000	/12 (255.240.0.0)	1,048,576 host addresses
	11111111.11111000.00000000.00000000	/13 (255.248.0.0)	524,288 host addresses
	11111111.11111100.00000000.00000000	/14 (255.252.0.0)	264,144 host addresses
	11111111.11111110.00000000.00000000	/15 (255.254.0.0)	131,072 host addresses
Class B	**11111111.11111111.00000000.00000000**	**/16 (255.255.0.0)**	**65,536 host addresses**
	11111111.11111111.10000000.00000000	/17 (255.255.128.0)	32,768 host addresses
	11111111.11111111.11000000.00000000	/18 (255.255.192.0)	16,384 host addresses
	11111111.11111111.11100000.00000000	/19 (255.255.224.0)	8,192 host addresses
	11111111.11111111.11110000.00000000	/20 (255.255.240.0)	4,096 host addresses
	11111111.11111111.11111000.00000000	/21 (255.255.248.0)	2,048 host addresses
	11111111.11111111.11111100.00000000	/22 (255.255.252.0)	1,024 host addresses
	11111111.11111111.11111110.00000000	/23 (255.255.254.0)	512 host addresses
Class C	**11111111.11111111.11111111.00000000**	**/24 (255.255.255.0)**	**256 host addresses**
	11111111.11111111.11111111.10000000	/25 (255.255.255.128)	128 host addresses
	11111111.11111111.11111111.11000000	/26 (255.255.255.192)	64 host addresses
	11111111.11111111.11111111.11100000	/27 (255.255.255.224)	32 host addresses
	11111111.11111111.11111111.11110000	/28 (255.255.255.240)	16 host addresses
	11111111.11111111.11111111.11111000	/29 (255.255.255.248)	8 host addresses
	11111111.11111111.11111111.11111100	/30 (255.255.255.252)	4 host addresses
	11111111.11111111.11111111.11111110	/31 (255.255.255.254)	2 host addresses
	11111111.11111111.11111111.11111111	/32 (255.255.255.255)	"Host Route"

Figure 1-2 *Classless IPv4 Address Allocation*

NAT with Private Addresses

NAT along with private IPv4 addresses has arguably been the reason IPv4 has survived all these years—to the deferment, and perhaps detriment, of IPv6. Although NAT was meant to be a "short-term" solution to the IPv4 exhaustion problem, for more than 20 years, it has been a critical tool used on almost every user network.

The following private IPv4 addresses are defined by RFC 1918:

- 10.0.0.0–10.255.255.255 (10.0.0.0/8)
- 172.16.0.0–172.31.255.255 (172.16.0.0/12)
- 192.168.0.0–192.168.255.255 (192.168.0.0/16)

You will most likely recognize the 192.168.0.0 or 10.0.0.0 address used by many home routers that perform NAT.

NAT allows multiple hosts using private IPv4 addresses within their internal network to share one or more public IPv4 addresses when accessing the global Internet. A public IPv4 address is one that is routable in the global Internet. In most cases, a network

connecting to the Internet needs only one or a few public IPv4 addresses. Internally, the network uses private IPv4 address space, providing up to 16 million unique addresses. Typically a single public IPv4 address can be used to translate several hundred or even a few thousand internal private IPv4 addresses. NAT was initially defined in RFC 1631, *The IP Network Address Translator (NAT)*.

Note Public IPv4 addresses are addresses that are globally routed by entities that provide Internet connectivity.

The Internet Assigned Numbers Authority (IANA) allocated three ranges of addresses that are considered private IPv4 addresses, as specified in RFC 1918. These private network addresses can be assigned to devices in private networks but are not routable in the global Internet. A private IPv4 address must be translated to a public IPv4 address before it can be forwarded by an Internet-facing router.

Note Although this is by far the most common use, NAT is not utilized only for translation between private and public IPv4 addresses. It can be used between any pair of addresses.

Note RFC 6598, *IANA-Reserved IPv4 Prefix for Shared Address Space*, assigns another block of shared address space, 100.64.0.0/10. Shared address space is distinct from RFC 1918 private address space because it is intended for use on service provider networks. However, it can be used in a manner similar to an RFC 1918 private address space.

Port address translation (PAT) is a form of dynamic NAT that maps multiple IPv4 addresses to a single IPv4 address using different TCP/UDP port numbers. PAT is also known as network address port translation (NAPT), NAT overloading, single-address NAT, and port-level multiplexed NAT. With PAT, each computer on a private network can be translated to the same public IPv4 address but with a different source port number assigned. Whenever the term NAT is mentioned in this book, it also includes PAT.

Figure 1-3 shows an example of NAT. In this case, each of the devices on the home network will receive an RFC 1918 private address—for example, a 192.168.1.0/24 address. The user's home router will perform NAT, translating all Internet-bound packets to a single public IPv4 address.

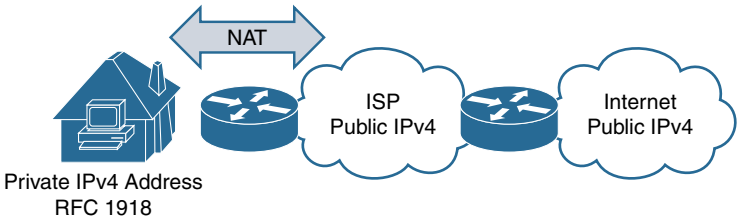

Figure 1-3 *Basic NAT for IPv4*

In many cases, the public IPv4 address is that of the router's Internet-facing interface. The home devices are usually clients, initiating communications to a server on the Internet with a public IPv4 address such as Facebook's or Amazon's. As the server responds to the client requests, the home router translates the public IPv4 address back to the client's original private IPv4 address.

Problems with NAT

In our need to conserve public IPv4 address space, we have tended to overlook some of the inherent problems with NAT. RFC 1631 even cautioned about negative effects of NAT: "It hides the identity of hosts. While this has the benefit of privacy, it is generally a negative effect." Although NAT was designed as a short-term solution, it became a rather permanent fixture in today's networks.

At the very least, NAT means that our routers, application gateways, firewalls, and other devices must perform extra processing to make NAT work, which also causes latency. The following are some of the major issues with NAT:

- **Checksum recalculations:** When carrying TCP segments, modifying the IPv4 header also means the IPv4 checksum must be recalculated.
- **ICMP manipulation:** Many ICMP messages, such as the Destination Unreachable message, embed in their payload the original IPv4 header that led to the ICMP message being generated. Because the original IPv4 address was translated by NAT, the NAT device must translate these addresses as well.
- **IPsec issues:** NAT invites complications when using Internet Protocol Security (IPsec) in transport mode. If IPsec AH (Authentication Header) is used, the NAT translation breaks the integrity check because the packet was modified during transport. NAT modifies the TCP/UDP checksum, which causes the integrity check to fail at the other end. Therefore, NAT cannot be used with IPsec in transport mode. Although it may work in tunnel mode, there can be issues in that case as well.
- **Breaking end-to-end reachability:** NAT causes problems with one of the key principles of the Internet: end-to-end reachability. For most users, this has not been an issue as the Internet has been comprised of mainly client/server networks. Generally, the Internet has been used by users and consumers who download content

such as web pages, email, and music. However, this is changing as the web continues to transition to more of a two-way medium, relying on video conferencing applications like Skype and smart home IoT applications that remotely manage the home. NAT makes accessing a device with a private IPv4 address difficult. Therefore, peer-to-peer, IoT, and many other types of services must provide an intermediary device, a kind of server with a public IPv4 address that both end devices can connect to—and that breaks pure end-to-end reachability. There are methods such as port forwarding to allow direct access to a device with a private IPv4 address, but they add another layer of complexity and potential problems.

- **Performance:** The process of having to translate addresses as packets leave and re-enter a network creates delay. This is especially a problem when NAT has to be performed more than once, such as when carrier-grade NAT (CGN) is used.

Many service providers use CGN, also called large-scale NAT (LSN), as another way to conserve IPv4 address space (see Figure 1-4). With CGN, instead of giving user networks their own public IPv4 addresses, a user's private IPv4 address is translated using NAT to another private IPv4 address to connect to its provider's network. The ISP typically uses an RFC 6598 address.

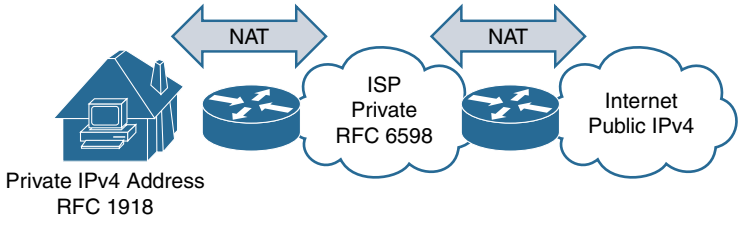

Figure 1-4 *Carrier-Grade NAT*

When the ISP forwards this packet to another provider (the public IPv4 Internet), NAT is performed once again, this time translating the RFC 6598 address to a public IPv4 address. The reverse process takes place for the returning packets. As you can imagine, this adds additional complexity and latency. For those familiar with port forwarding, CGN makes this even more difficult. (Port forwarding is the process of intercepting data traffic and modifying a combination of the IP address and TCP/UDP port number in order to redirect the data to a specific device.)

NAT is Not Security

Many people have argued and continue to argue the security merits of NAT. NAT has certain side effects that resemble security, but NAT itself is not a security feature. The IETF never intended NAT to be used for security. Many have also argued and demonstrated that NAT actually makes a network less secure and introduces its own security problems. RFC 1631 states that "NAT reduces the number of options for providing security" and makes debugging more difficult. Another excellent reference is RFC 2993, *Architectural Implications of NAT*.

However, we can't ignore the fact that one of the *perceived* security benefits of NAT is that it hides internal, private IPv4 addresses from the outside world. This makes accessing a private network address from the public Internet very difficult. Without getting too deep into this topic, it really isn't NAT that provides this apparent security. What provides the security is that NAT requires *state*. A NAT device must remember which private addresses and port numbers got mapped to which public addresses and port numbers. This statefulness means that any device performing NAT must also act as a stateful firewall.

I think we can just about all agree that what provides security for a network is a stateful firewall or other type of security appliance. After all, if NAT really did provide the security we deem necessary for our networks, we wouldn't have any issues with viruses, malware, ransomware, or any of the other problems plaguing our devices. Besides, most devices hiding behind NAT get infected by malicious email attachments and their users end up going to rogue websites, downloading infected files, falling for phishing schemes, and even using fake antivirus software.

For more information about the perceived security nature of NAT, you may want to check out these articles and this excellent and very funny video:

- **"Is NAT a Security Feature?" by Ivan Pepelnjak:** blog.ipspace.net/2011/12/is-natsecurity-%20feature.html
- **"IPv6 Security Myth #3—No IPv6 NAT Means Less Security" by Chris Grundemann:** www.internetsociety.org/deploy360/blog/2015/01/ipv6-security-myth-3-no-ipv6-nat-means-less-security/
- **"'Fanboy' Series—IPv6 and NATs" by Andrew Yourtchenko:** www.youtube.com/watch?v=v26BAlfWBm8

Is there NAT for IPv6? Yes, but it's not the same as NAT for IPv4. Most of the versions of NAT for IPv6 have to do with protocol translation between IPv6 and IPv4. NAT66 (described in a long-expired draft RFC; see tools.ietf.org/html/draft-mrw-behave-nat66-02) and NPTv6 (described in RFC 6296) are used for translation between two IPv6 addresses. NAT66 and NPTv6 are both designed to support address independence, not stateful translation between public and private IPv6 addresses. The various IPv6 NAT options are explained further in Chapter 17, "Deploying IPv6 in the Network."

NAT Example

Figure 1-5 shows an example of a network that uses NAT. Company XYZ has the public IPv4 address 209.165.200.248/29. This gives the company only six usable host addresses: 209.165.200.249 through 209.165.200.254. Company XYZ has hundreds of hosts, all of which might need access to the Internet at any given time. Because the company has only six public IPv4 addresses, it uses private IPv4 address space and NAT to meet its Internet access needs.

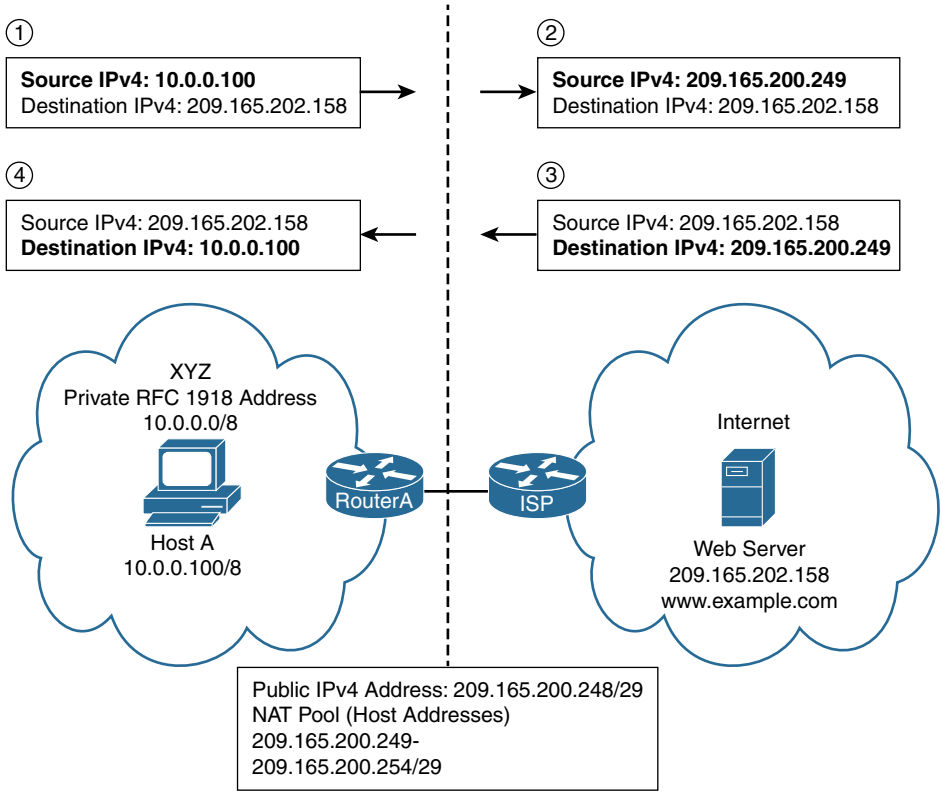

Figure 1-5 *Example of IPv4 NAT*

Company XYZ's network administrator has chosen to use the private network address 10.0.0.0/8. This provides more than enough host IPv4 addresses to meet the company's needs. RouterA is Company XYZ's router, which connects the private network to the Internet by forwarding packets to the ISP router. The steps in Figure 1-5 show how NAT translates a private IPv4 address to and from a public IPv4 address. For simplification, the port translations that are part of PAT have not been included.

Step 1. HostA on Company XYZ's network sends a packet with its source IPv4 address 10.0.0.100 and the destination IPv4 address 209.165.202.158 to the www.example.com web server.

Step 2. RouterA receives the packet and performs the NAT translation before forwarding the packet to the ISP router. RouterA replaces the packet's private source IPv4 address of 10.0.0.100 with the public address 209.165.200.249. The router also likely change the source TCP/UDP port number as part of NAT. (Again, PAT has not been included in this discussion to keep it simple.)

RouterA maintains a NAT table to keep track of all the addresses and port assignments that it modifies so that it can put the original address and port numbers back for the return traffic in step 4.

Step 3. The www.example.com web server returns the packet to the host with the source IPv4 address 209.165.202.158 and the destination IPv4 address 209.165.200.249.

Step 4. RouterA receives the return traffic for HostA. Using its NAT table, RouterA identifies the destination IPv4 address and destination port number, and it modifies these with the original address, 10.0.0.100, and the original port number.

What About IPv5?

Whatever happened to IPv5? No, the Internet doesn't have an issue with odd or prime numbers. In the late 1970s, a family of experimental protocols, known as Internet Stream Protocol (ST) and later ST2, was developed. Originally defined in IEN-119 (1979), ST was later revised in RFC 1190 and RFC 1819.

ST was an experimental resource reservation protocol intended to provide Quality of Service (QoS) for real-time multimedia applications such as video and voice. ST consisted of two protocols—ST (Internet Stream Protocol) and Stream Control Message Protocol (SCMP). Internet Stream Protocol version 2 (ST-II or ST2) was not designed as a replacement for IPv4. The idea was that a multimedia application would use both protocols—IPv4 for the transfer of traditional packets and ST2 for packets carrying real-time data.

Although it was never recognized as IPv5, when encapsulated in IPv4, ST uses IP Protocol Number 5 (RFC 1700). In other words, although it was never implemented, the designation "IP version 5" was already taken. Today's standard for resource reservation is the transport layer protocol Resource Reservation Protocol (RSVP), which can be used to provide receiver-initiated setup over IPv4. RSVP is described in RFC 2205.

The Fascinating History of IPv6

The story of how IPv6 came to be the protocol it is today is a fascinating one. Really! We might just assume that protocols are always developed with a specific logical process similar to the scientific method, with the final result being the one-and-only inevitable solution. Well, that's not always how it works, as is the case with IPv6.

The story of how IPv6 became the successor to IPv4 is a story that you might not have expected. It illustrates both the technical and human sides of protocol development—a process that at times brings intense disagreement on what technology or standard is best. The story behind IPv6 also shows how Internet governance went from being controlled mostly by individuals from American companies, universities, and research institutions to a more international collaboration.

Some Background

To get a bit of background on the development of IPv6, let's look at the world of the Internet back in 1990:

- In August 1990 the first web page was accessible.
- There were about 3 million users, and 73% of them were in the United States.
- There was no Amazon, no Google, no Yahoo, no Facebook. (There was not even Netscape, which you may not be old enough to remember.)
- The most popular Internet application was text-based email.

It's important to note that up to this point, the Internet was very US-centric. Not only were most of the users Americans, but the governing organizations such as the Internet Architecture Board (IAB; originally called the Internet Configuration Control Board) and the IETF consisted primarily of Americans working at US corporations, universities, and research institutions.

The Advanced Research Projects Agency Network (ARPANET), the precursor to today's Internet, was developed by the US Department of Defense. ARPANET came to life in 1969, interconnecting the University of California Los Angeles (UCLA), the Augmentation Research Center at Stanford Research Institute (SRI), the University of California Santa Barbara (UCSB), and the University of Utah. And for the next 20 years, the Internet was primarily controlled by Americans at US institutions.

Note To learn more about the creation of the Internet and those who were responsible for its development, I highly recommend the national bestselling book, *Where Wizards Stay Up Late: The Origins of the Internet*, by Katie Lyons.

Established in 1979 by Vint Cerf, the IAB is responsible for the direction of the Internet, Internet protocol architecture, and the ultimate responsibility for approving protocols. In 1990 the IAB consisted of 11 members (there are 13 today), mostly Americans with a core philosophy of "working code and rough consensus."

The IAB established the IETF in 1986. The IETF had no formal membership, all were volunteers—again mostly Americans. Its primary role was and still is to develop and propose standards (protocols) that make up the Internet protocol suite, TCP/IP.

In 1990 it was not a foregone conclusion that TCP/IP would be the Internet protocol suite of the future. The International Organization for Standardization (ISO) was ready with its own protocol suite, the Open Systems Interconnection (OSI) that it had been developing with the International Telecommunications Union (ITU). ISO, founded in 1947, in Geneva, Switzerland, is an international organization that promotes worldwide commercial and industrial standards. TCP/IP had become the de facto standard (which means it was the standard only because it was so widely used). But ISO still thought its OSI suite would ultimately become the officially endorsed standard of the Internet

(the de jure standard). It was just a matter of time before these two standards organizations would face off, and it would happen over IPv6.

Note ISO is not an acronym but refers to the Greek word *isos*, which means "equal."

IPv4 Address Exhaustion and the Need for More International Involvement

In 1990, both the IAB and the IETF began discussions about the shortage of IPv4 addresses. The sizes of Internet routing tables were increasing rapidly, and the number of Internet users was exploding. The IAB and the IETF agreed that it was time to begin designing and testing a new network layer protocol as the successor to IPv4. In addition to increasing the size of the address space, this was also a unique opportunity to fix the limitations of IPv4 and develop a protocol to ensure reliable growth and enhanced performance for the future.

Note Much of this process is well documented in RFC 1752, *Recommendation for IPng*.

During an IETF meeting in August 1990, it was estimated that current IPv4 address assignment rate would be depleted by 1994. This was even before the arrival of the World Wide Web, which was soon to revolutionize how we use the Internet. At the time, the Web was still under development by Tim Berners-Lee at CERN (French acronym for the European Organization for Nuclear Research in Switzerland). As we would soon see, this would bring a vast number of new Internet users that would greatly add to the problem IPv4 address depletion.

At this time, the IAB also had concerns about the Internet being so US-centric. Not only were the IAB, IETF, and other Internet governing bodies made up of mostly Americans, but the US had most of the IPv4 address space. At the time, IPv4 addresses were allocated by a single US entity, IANA, and the US had received the majority of IPv4 addresses.

Note In 1990, the process of IPv4 address allocation was delegated by IANA to SRI International's Networking Information Center in Marina Del Rey, California, which was funded by the US Department of Defense.

Vint Cerf, IAB chair, wanted to see more international participation. He advocated that IPv4 addresses should be allocated by international organizations. This is how the five RIRs were eventually created.

The battle for the Internet suite of protocols had two contenders: IETF's TCP/IP suite and ISO's OSI suite. TCP/IP was the dominant suite, with a working set of protocols. TCP/IP

also had the backing of the IETF, which documented all protocols and processes using RFCs (Request for Comments). OSI had limited deployment, and the management of its standards used a much more cumbersome and bureaucratic process (at least in the eyes of the IETF). OSI also had the backing of the US National Institute of Standards and Technology (NIST) and large US corporations like Digital Equipment Corporation (DEC).

And the winner is…. In the summer of 1992, the IAB had a meeting in Kobe, Japan. Without consulting with the IETF, as was customary, the IAB, in an effort to be more internationally inclusive, proposed OSI's Connectionless-mode Network Protocol (CLNP) as the replacement for IPv4. The CLNP proposal was called TCP and UDP with Bigger Addresses (TUBA). It was to replace IPv4 and its 32-bit address space with a variable-length, 160-bit CLNP address.

And the name of the new protocol? The protocol was to be named IPv7, erroneously skipping IPv6. What about the TCP/IP protocol suite? The plan was to use as much of the upper layers of TCP/IP as possible, which is why TUBA included "TCP and UDP" in its name.

When the IAB presented this proposal to the IETF, there was immense outrage and disbelief among the IETF members. First of all, they were furious that this was done without consulting them, not to mention discarding IPv4 in its entirety. The IETF also feared losing control over the standards development process to the more complex and bureaucratic ISO. The IETF's message was blunt and unified: Reject the IAB's mandate to replace IPv4 with CLNP.

A Call for Proposals

Vint Cerf and the IAB got the message. In 1993, IETF decided to open up the process with a call for proposals to replace IPv4. The announcement was made in RFC 1550, *IP: Next Generation (IPng) White Paper Solicitation*. This is why you might see some older documentation refer to IPv6 as IPng.

Note RIPng is Routing Information Protocol next generation, which is RIP for IPv6.

The Internet Engineering Steering Group (IESG), headed by Scott Bradner and Allison Mankin, would make the final decision. Three proposals were submitted:

- **Common Architecture for the Internet (CATNIP):** CATNIP proposed integrating IPv4, Internetwork Packet Exchange (IPX), and CLNP. IPX was part of the Internetwork Packet Exchange/Sequenced Packet Exchange (IPX/SPX) suite of protocols, used primarily on networks employing the Novell NetWare operating systems. CLNP is an OSI standard defined in ISO 8473 and is the equivalent of IPv4 for the OSI suite of protocols. CATNIP is defined in RFC 1707.

- **Simple Internet Protocol Plus (SIPP):** SIPP recommended increasing the IPv4 address size from 32 bits to 64 bits, along with making additional improvements to the IPv4 header for more efficient forwarding. SIPP is defined in RFC 1710.
- **TCP/UDP over CLNP-Addressed Networks (TUBA):** TUBA was aimed at minimizing the risk associated with migration to a new IP address by replacing IPv4 with CLNP and its address size of 20 bytes (160 bits). TCP, UDP, and the traditional TCP/IP applications would run on top of CLNP. TUBA is defined in RFCs 1347, 1526, and 1561.

The IETF formed the IP Next Generation working group in 1993. In July 1994, it was no surprise that the IETF announced its recommendation that the IP-based SIPP, written by Steve Deering, Paul Francis, and Bob Hinden, but with an address size of 128 bits, should replace IPv4.

A More IP Version of IPv6

In 1994, the IESG chose SIPP, but with a 128-bit IP address. The choice of SIPP with a 128-bit IP address meant that the IETF would continue to retain control of the standards development process.

In 1995, IETF published RFC 1883, *Internet Protocol, Version 6 (IPv6) Specification*, which later became obsolete and was replaced by RFC 2460 in 1998. As mentioned earlier, the name IPv5 was already taken, and the next available version was IPv6. So, in 2001, the IPng working group was renamed the IPv6 working group.

In 1996, the IETF formed an IPv6 test bed network known as *6bone*. The 6bone network started out using IPv6 over IPv4 tunneling/encapsulation, operating over the IPv4-only Internet to support IPv6 transport. Later it migrated to native links for IPv6. 6bone was phased out in 2006.

RIRs began allocating IPv6 addresses to their customers in 1999. Adoption was slow at first, but in 2007, RIRs began receiving significantly more requests for IPv6 address space. This was because RIRs supported a more widespread deployment of IPv6. Recently, with IANA running out of IPv4 addresses and now with four out of five RIRs depleted of IPv4 addresses, IPv6 adoption is seeing much more rapid growth.

In 2000, many vendors began adding IPv6 support to their mainstream products. Cisco Systems added IPv6 support with Cisco IOS Software Release 12.2(2)T. Linux vendors started supporting IPv6 in 2000, and Microsoft announced IPv6 support with Windows XP in 2001.

The IAB, IETF, RIRs, and other Internet governing bodies are much more international today than they were in 1990. The desire of the IAB to be a more inclusive collaboration of representatives from countries all over the world has become more of a reality.

Note Besides RFC 2460, there are dozens of other RFCs related to IPv6. An excellent list of these RFCs can be found at sites.google.com/site/ipv6center/ipv6-rfcs.

Note To learn more about the politics and personal side of protocol development, including IPv6, I highly recommend the book *Protocol Politics: The Globalization of Internet Governance (Information Revolution and Global Politics)*, by Laura DeNardis.

IPv6: More Than Just Longer Addresses

Many people who are new to IPv6 are surprised to learn that IPv6 is much more than just a larger source and destination IP address. Developers of IPv6 took this opportunity to not only improve IP but also many of the protocols and processes related to IP. This foresight has continued, with standards and informational RFCs that go well beyond the Internet of today. For example, the home of tomorrow may look more like an enterprise network of today. Perhaps a home will be connected to multiple ISPs, each providing differentiated services. A home network may include multiple subnets for personal use, business needs, entertainment systems, guest access, smart home devices, and so on. A good discussion of these home networking concepts can be found in RFC 7268, *IPv6 Home Networking Architecture Principles*.

So, why an entire book on IPv6? As mentioned earlier, IPv6 is much more than just a larger IP address. The following is just a sampling of some of the important features of IPv6 that are discussed in this book:

- An IPv6 address (global unicast address) includes a Subnet ID field that can be defined by the organization so there is no need to borrow bits from the host portion of the address to create subnets. This allows the organization to easily design and manage subnets for its own purposes, not for conserving address space as with IPv4.
- The IPv6 header has been significantly modified to include one new field and the removal of others.
- A new ICMPv6 Neighbor Discovery Protocol (NDP) has four new message types. ICMPv6 is used instead of ARP for resolving Layer 2 to Layer 3 addresses. NDP is also used by IPv6 routers to send messages suggesting how a device can dynamically receive its IPv6 addressing information.
- A non-routable IPv6 link-local address has significant uses in IPv6, including during neighbor discovery.
- A new solicited-node multicast address is used instead of a broadcast address to make neighbor discovery more efficient.

- IPv6 includes Stateless Address Autoconfiguration (SLAAC), which allows devices to obtain unique and globally routable addresses without the services of a DHCPv6 server.
- Stateless and stateful DHCPv6 provides several options for a device to obtain some or all of its addressing information.
- There is a neighbor cache, which is similar to the ARP cache for IPv4 but with multiple states.
- Devices can have multiple IPv6 addresses on the same subnet or different subnets to allow for greater privacy and additional features. These options can be managed at the host or network level.

And there is much more as well. Don't worry about understanding these concepts right now. This book explains them and many more of the nuances of IPv6 in a step-by-step, simple manner.

IPv6 Myths

There are several misperceptions or myths regarding IPv6. IPv6 has been around quite some time, initially introduced in 1995 with RFC 1883 and later obsoleted with RFC 2460 in 1998. Over the years, as IPv6 evolved and as people discussed the merits of the new protocol, certain misconceptions ensued. Let's take a look at some of them:

- **IPv6 is more secure than IPv4.** Some may think that IPv6 is more secure than IPv4 because it uses IPsec. IPsec is a security protocol suite for authenticating and encrypting IP packets that was initially developed for IPv4. IPsec does not make IPv6 any more secure than IPv4. IPsec can't stop all attacks against IPv6, ICMPv6, or any of the other related protocols. The fact is that neither IPv4 nor IPv6 is any more secure than the other. IPsec was originally required on all implementations of IPv6, but RFC 6434 changed it to be only a recommendation.
- **IPv6 is less secure than IPv4.** Some believe that IPv6 is less secure than IPv4 because it doesn't use NAT. As we have already discussed, NAT is not security. Although NAT provides a side effect of hiding your private IPv4 addresses from the public Internet, this does not equate to security.
- **IPv6 will replace IPv4.** There is no switchover date for going from IPv4 to IPv6. IPv4 and IPv6 will coexist for the foreseeable future. So don't worry; you still get to use all of your IPv4 knowledge along with your newfound IPv6 knowledge.
- **IPv6 isn't necessary.** Some believe that because they have plenty of IPv4 addresses, they don't need to understand or implement IPv6. As mentioned at the beginning of this chapter, for both business and security reasons, it is important to begin to implement IPv6. Ignoring IPv6 could potentially isolate your online services from some consumers, and it could also make your network vulnerable to certain types of IPv6 MITM and DoS attacks.

- **IPv6 is too complex.** IPv6 and related protocols such as ICMPv6 may be new and different, but they are not any more complex than IPv4. In many ways, IPv6 is actually much easier. Thanks to the Subnet ID field in the IPv6 header, subnetting IPv6 is simple. If you can count in hexadecimal, then you can subnet in IPv6. In addition, there is no fragmentation of IPv6 packets by routers, so these fields don't exist in the IPv6 header. IPv4's fragmentation and reassembly can be a complicated process to understand.
- **IPv6 improves QoS.** Another common misconception is that IPv6 provides better QoS. Both IPv4 and IPv6 use the same Differentiated Serves and Integrated Services fields for QoS. IPv6 does provide a new Flow Label field that has the potential to improve the efficiency of flows in an IPv6 network. Although many operating systems set the Flow Label field for IPv6 packets, currently there aren't a lot of implementations that look at it. The Flow Label field is discussed in Chapter 3, "Comparing IPv4 and IPv6."

Transitioning to IPv6

> *"The IPv6 protocol standard is critical for supporting the Internet's continued development. Network operators, content providers, software and hardware developers, and enterprises, among others, need to implement IPv6 in order to ensure efficiency, global connectivity, and long-term growth of the Internet."*
>
> —Internet Society (ISOC) Policy Brief: Adoption of IPv6, April 2016

When will the world transition to IPv6? Well, there is no specific date to switch from IPv4 to IPv6. IPv4 and IPv6 will most likely coexist for many years. The transition to IPv6 is in progress and will continue to gain momentum for the reasons we have discussed in this chapter. In June 2011, ISOC held its first World IPv6 Day. The goal was for service and content providers to test their IPv6 implementations. Since then, IPv6 adoption by service and content providers has continued to increase, as has the number of networks enabled for IPv6.

IPv6 has a variety of tools to help with the transition from IPv4 to IPv6, including tunneling and NAT. Tunneling encapsulates an IPv6 packet into an IPv4 packet so that it can be delivered over IPv4-only networks. NAT provides a mechanism to translate an IPv4 address to an IPv6 address or an IPv6 address to an IPv4 address. However, these are just temporary transition tools and in most cases no longer needed.

Whether with a home network or a large enterprise network, the goal should be native IPv6 connectivity. Native IPv6 means no translation techniques or tunneling protocols are needed. Large service providers such as Comcast, AT&T, Time Warner, Verizon, Sky Broadband, Deutsche Telekom, and many others are offering native IPv6. If your ISP is not providing native IPv6 connectivity, ask when it plans to deliver IPv6.

As we mentioned earlier, the killer application for IPv6 is the Internet itself. As ISOC, IETF, and the RIRs have been saying for years, the Internet cannot sustain itself with IPv4. It can't continue to grow and evolve without more addresses, and the only long-term solution is IPv6.

Now is the best time for IT departments to begin preparation for IPv6—before there is a pressing need or an immediate requirement. IT organizations need time to familiarize themselves with IPv6, educate their workforces, and test their IPv6 implementations. Preparation should focus on all aspects of the IT department, such as network operations, security, application developers, content developers, and the data center, just to name a few.

Here are some of the things that can be done to start the migration to IPv6:

- **Include IPv6 in your IT strategy:** IPv6 needs to be a part of your IT strategic planning for both the short term and the long term. IPv6 needs to be integrated into your IT planning processes, status meetings, training, and purchasing plans.
- **Provide training:** The optimum time to begin training about IPv6 or at least becoming familiar with it is before you have a requirement to implement it. Education/training is the best way for IT staff to lose their fear and intimidation of IPv6. Training can also include non-IT staff, such as salespeople. A friend of mine who's been a network engineer for over 25 years told me he was hoping to have retired by the time he had to learn IPv6. As you read this book, you will see that IPv6 is not as complicated as you might think; in many ways, it's easier to deal with than IPv4. For example, subnetting doesn't require any bitmapping, so IPv6 subnets can be easier to configure and identify.
- **Inventory equipment, applications, and services:** Develop a list of devices on your network (or use an existing inventory list) and check whether these devices currently support IPv6 or what upgrade/replacement is necessary. Include the following:
  - Clients and servers
  - Routers and multilayer switches
  - Layer 2 switches
  - Firewalls and other security devices
  - Printers, webcams, phones, and any other devices that have IP addresses on your network
  - Applications, operating systems, and software services
- **Keep an eye on purchasing:** Your purchasing plan should confirm that any new equipment purchased is IPv6-compatible. Ask vendors whether their hardware/software supports IPv6. If it doesn't, ask them what their product implementation plan is for including IPv6 support.
- **Create a test lab:** Creating a lab that is off the production network is a great way to learn and play with IPv6. Then, as you become more comfortable with IPv6, you can

introduce it into your network as you continue to run IPv4. The lab should simulate the operational network in all aspects. It should come as close to the live network environment as possible—with the same routers, switches, firewalls, load balancers, applications, end systems, and so on.

Summary

After reading this chapter, you should have a better understanding of the limitations of IPv4 and why it is important to begin transitioning to IPv6.

Here are some of the key items discussed in this chapter:

- How we use the Internet today is much different than how it was used when IPv4 was developed. Today we have more users, more devices, and new demands. We have moved from just an *Internet of computers* to also an Internet of things.
- Although no one knows exactly when, we will eventually run out of IPv4's 4.29 billion addresses, but the fact is that the Internet is in the final stages of public IPv4 address availability.
- IPv6, with its 128-bit address scheme, provides more than enough globally unique addresses to support the growth of the Internet.
- IPv4 and IPv6 will coexist for the foreseeable future. IPv6 includes tools and migration strategies that allow both protocols to coexist.
- The combination of CIDR, NAT, and private addressing has helped slow the depletion of IPv4 address space. However, NAT introduces complexity, latency, and other problems.
- One of the perceived security benefits of NAT is that it provides address hiding. But NAT is not considered security. Because NAT requires state, a NAT device must act as a stateful firewall.
- In addition to providing a larger address space, IPv6 offers additional enhancements, such as ICMPv6 Neighbor Discovery Protocol (NDP) for Layer 3 to Layer 2 address resolution, and Stateless Address Autoconfiguration (SLAAC) that allows a device to create an IPv6 address without the services of DHCP.
- Now is the best time for IT departments to begin familiarizing themselves with IPv6.
- The Internet is the killer application for IPv6. IPv6 is the successor to IPv4 that will ensure the growth of the Internet.

Review Questions

1. What are two reasons network operators need to begin to transition to IPv6?
2. What temporary measures did the IAB and IETF implement when they recognized the impending shortage of IPv4 addresses?

3. List two problems with NAT.
4. Is NAT considered security? Why or why not?
5. Why isn't IPv5 the replacement for IPv4?
6. In 1992, what protocol did the IAB recommend as the replacement for IPv4?
7. In 1994, what protocol did the IESG recommend and the IAB select as the replacement for IPv4?
8. List three myths regarding IPv6.

References

Endnotes

1. Chetan Sharma Consulting, *US Wireless Market Update Q1 2016*, www.chetansharma.com/usmarketupdateq12016.htm.
2. *World IPv6 Launch Measurements*, www.worldipv6launch.org/measurements/.
3. *Case Study: T-Mobile US Goes IPv6-Only Using 464XLAT*, www.internetsociety.org/deploy360/resources/case-study-t-mobile-us-goes-ipv6-only-using-464xlat/.
4. APNIC, *IPv6 Performance—Revisited*, blog.apnic.net/2016/08/22/ipv6-performance-revisited/.
5. Internet Society, *Facebook News Feeds Load 20-40% Faster over IPv6*, www.internetsociety.org/deploy360/blog/2015/04/facebook-news-feeds-load-20-40-faster-over-ipv6/.
6. GlobalWebIndex, *Digital Consumers Own 3.64 Connected Devices*, www.globalwebindex.net/blog/digital-consumers-own-3.64-connected-devices.
7. *IPv4 Address Report*, www.potaroo.net/tools/ipv4.
8. *APRINIC IPv6 Programme*, afrinic.net/services/ipv6-programme.

RFCs

RFC 760, *DoD Standard, Internet Protocol*, USC, IETF, www.ietf.org/rfc/rfc760.txt, January 1980.

RFC 791, *Internet Protocol, DARPA Internet Program Protocol Specification*, USC, www.ietf.org/rfc/rfc791.txt, September 1981.

RFC 1190, *Experimental Internet Stream Protocol, Version 2 (ST-II)*, CIP Working Group, www.ietf.org/rfc/rfc1190.txt, October 1990.

RFC 1338, *Supernetting: An Address Assignment and Aggregation Strategy*, V. Fuller, BARRNet, www.ietf.org/rfc/rfc1338, June 1992.

RFC 1347, *TCP and UDP with Bigger Addresses (TUBA), A Simple Proposal for Internet Addressing and Routing*, Ross Callon, DEC, www.ietf.org/rfc/rfc1347.txt, June 1992.

RFC 1519, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*, V. Fuller, BARRNet, www.ietf.org/rfc/rfc1519.txt, September 1993.

RFC 1526, *Assignment of System Identifiers for TUBA/CLNP Hosts*, D. Piscitello, Bellcore, www.ietf.org/rfc/rfc1526.txt, September 1993.

RFC 1550, *IP: Next Generation (IPng) White Paper Solicitation*, S. Bradner, Harvard University, www.ietf.org/rfc/rfc1550.txt, December 1993.

RFC 1561, *Use of ISO CLNP in TUBA Environments*, D. Piscitello, Core Competence, www.ietf.org/rfc/rfc1561.txt, December 1993.

RFC 1700, *Assigned Numbers*, J. Reynolds, ISI, IETF, www.ietf.org/rfc/rfc1700.txt, October 1994.

RFC 1707, *CATNIP: Common Architecture for the Internet*, M. McGovern, Sunspot Graphics, www.ietf.org/rfc/rfc1707.txt, October 1994.

RFC 1710, *Simple Internet Protocol Plus White Paper*, R. Hinden, Sun Microsystems, www.ietf.org/rfc/rfc1710.txt, October 1994.

RFC 1752, *The Recommendation for the IP Next Generation Protocol*, S. Bradner, Harvard University, www.ietf.org/rfc/rfc1752.txt, January 1995.

RFC 1819, *Internet Stream Protocol Version 2 (ST2) Protocol Specification - Version ST2+*, ST2 Working Group, www.ietf.org/rfc/rfc1819.txt, August 1995.

RFC 1883, *Internet Protocol, Version 6 (IPv6) Specification*, S. Deering, Xerox PARC, www.ietf.org/rfc/rfc1883.txt, December 1995.

RFC 1918, *Address Allocation for Private Internets*, Y. Rekhter, Cisco Systems, IETF, www.ietf.org/rfc/rfc1918.txt, February 1996.

RFC 2205, *Resource ReSerVation Protocol (RSVP)—Version 1 Functional Specification*, R. Braden, ISI, www.ietf.org/rfc/rfc2205.txt, September 1997.

RFC 2235, *Hobbes' Internet Timeline*, R. Zakon, MITRE, www.ietf.org/rfc/rfc2235.txt, November 1997.

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*, S. Deering, Cisco Systems, www.ietf.org/rfc/rfc2460.txt, December 1998.

RFC 2993, *Architectural Implications of NAT*, T. Hain, Microsoft, www.ietf.org/rfc/rfc2993.txt, November 2000.

RFC 3330, *Special-Use IPv4 Addresses*, IANA, www.ietf.org/rfc/rfc3330.txt, September 2002.

RFC 3927, *Dynamic Configuration of IPv4 Link-Local Addresses*, S. Cheshire, Apple Computer, www.ietf.org/rfc/rfc3927.txt, May 2005.

Websites

Google's list of IPv6 RFCs, sites.google.com/site/ipv6center/ipv6-rfcs

Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2015–2020 White Paper, www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/mobile-white-paper-c11-520862.html

Internet World Stats, www.internetworldstats.com/stats.htm

North American Network Operators' Group (NANOG) IPv6 presentations, Search for IPv6, www.nanog.org/archives/presentations

This page intentionally left blank

Chapter 2

IPv6 Primer

"It ain't an easy job but when you bring a herd into town and you ain't lost one of them, ain't a feeling like it in the world."

—Cat Herders, Electronic Data Systems (EDS) commercial, aired during Super Bowl XXXIV

Learning IPv6 is a little like herding cats. Where to start? There are a lot of new concepts, and wherever you start, it always seems like it might have been better to have started somewhere else. This is because each topic relies on some previous knowledge of some of the other topics.

This chapter provides a brief introduction to concepts, protocols, and terminology that will give you the information you need to help you better understand the rest of the book. This chapter also gives you a better idea of some of the changes to IPv6 and related protocols such as ICMPv6 (Internet Control Message Protocol version 6).

Except for the first topic on the hexadecimal number system, everything else discussed in this chapter will be re-introduced and discussed more thoroughly in later chapters. So don't worry about completely understanding any of these IPv6 protocols and concepts for now. After reading this chapter, you will probably have more questions. But rest assured: They will be addressed later in this book.

This chapter introduces the following topics:

- **Hexadecimal number system:** IPv6 uses hexadecimal numbers to represent IPv6 addresses. This section is intended for those who are new to the hexadecimal number system or need a little review.
- **IPv6 address types:** IPv6 addresses include new address types and changes to some of the address types used in IPv4. This chapter looks at four of the most common address types: global unicast, link-local unicast, unspecified unicast, and

solicited-node multicast. Understanding these types will help you better understand many of the protocols and message types that use these different types of addresses, such as ICMPv6.

- **Address terminology:** IPv6 uses some terms that may be unfamiliar, such as *prefix* and *prefix length*. This chapter explains some of these basic terms that are used throughout the book.
- **ICMPv6 Neighbor Discovery Protocol (NDP):** ICMPv6 is similar to ICMP for IPv4. However, ICMPv6 has been extended, via ICMPv6 NDP, to support additional functions. ICMPv6 has four new message types: Neighbor Solicitation (NS), Neighbor Advertisement (NA), Router Solicitation (RS), and Router Advertisement (RA). These messages are used for address resolution (similar to Address Resolution Protocol [ARP] in IPv4) and dynamic address allocation. Getting a brief look at these messages now will help you better understand IPv6 address types and dynamic address allocation.
- **Dynamic address allocation:** IPv6 devices can receive their addressing information without the services of a Dynamic Host Configuration Protocol version 6 (DHCPv6) server. In this chapter we will introduce the three methods sent by a router's ICMPv6 Router Advertisement message for dynamic address allocation.

Again, everything we introduce in this chapter is also discussed in much more detail in later chapters.

Hexadecimal Number System

An IPv6 address is 128 bits in length, and, as we will see, hexadecimal is the ideal number system for representing long strings of bits. An IPv6 address with long strings of hexadecimal digits may look intimidating at first, but don't worry. Chapter 4, "IPv6 Address Representation and Address Types," looks at how IPv6 addresses are represented. You will see that in most cases, hexadecimal numbers make it easier to determine the different parts of an address.

If you understand the decimal, or base 10, number system, you can understand any number system, including the hexadecimal, or base 16, number system. You may already be familiar with binary, or base 2, but even if you're not, you will still be able to understand base 16. The same general rules apply to all number systems.

When looking at integer-based number systems, there are three general rules:

- **Rule #1:** Base *n* number systems have *n* number of digits:
  - The base 10 (decimal) number system has 10 digits.
  - The base 2 (binary) number system has 2 digits.
  - The base 16 (hexadecimal) number system has 16 digits.

- **Rule #2:** All number systems begin with 0.

  Combining Rule #1 and Rule #2, we get:

  - Base 10 has 10 digits, starting with 0: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9
  - Base 2 has 2 digits, starting with 0: 0, 1
  - Base 16 has 16 digits, starting with 0: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, A, B, C, D, E, F

- **Rule #3:** The rightmost column, or least significant digit, is always the column of 1s (ones). Each preceding column is *n* times the previous column (where *n* is the base *n* number system). Using base 10 as an example, the first column is the column of 1s; the next column is 10 times the 1s column, or a column of 10s; the next column is 10 times the 10s column, or a column of 100s; and so on. Understanding this makes it very easy to convert other number systems to base 10 (see Table 2-1).

Table 2-1 *Number Systems*

Base *n* Number System	n^3	n^2	n^1	n^0
Base 10	1,000	100	10	1
Base 2	8	4	2	1
Base 16	4,096	256	16	1

- As you can see in Table 2-1, the columns in the three number systems include:
  - **Base 10:** 10,000s, 1000s, 100s, 10s, 1s
  - **Base 2:** 128s, 64s, 32s, 16s, 8s, 4s, 2s, 1s
  - **Base 16:** 4,096s, 256s, 16s, 1s

When you understand these three rules, you're ready to examine the hexadecimal number system more closely. The hexadecimal number system, base 16, has 16 digits, beginning with 0. Table 2-2 shows these 16 digits and their equivalents in decimal and binary.

Table 2-2 *Decimal, Hexadecimal, and Binary*

Decimal (Base 10)	Hexadecimal (Base 16)	Binary (Base 2)
0	0	0000
1	1	0001
2	2	0010
3	3	0011
4	4	0100
5	5	0101
6	6	0110
7	7	0111
8	8	1000
9	9	1001
10	A	1010
11	B	1011
12	C	1100
13	D	1101
14	E	1110
15	F	1111

Let's apply the three rules to the hexadecimal number system:

- **Rule #1:** The hexadecimal number system has 16 digits.
- **Rule #2:** Table 2-2 illustrates the 16 hexadecimal digits and their decimal and binary equivalents starting with 0. Notice that you needed unique alphanumeric digits, A through F, to represent the decimal values 10 through 15.
- **Rule #3:** For representing IPv6 addresses in hexadecimal, you only need to use the rightmost column, or the column of 1s.

Using Table 2-2, it is easy to convert from one number system to any of the others. With one known value, you can easily determine the values for the other two number systems. The following are three examples, in which subscripts are used to differentiate between the number systems:

- If your known value is 14_{10}, you can find that it equates to E_{16} and 1110_2.
- If your known value is A_{16}, you can find that it equates to 10_{10} and 1010_2.
- If your known value is 0010_2, you can find that it equates to 2_{10} and 2_{16}.

Why use hexadecimal numbers to represent IPv6 addresses? Hexadecimal is a natural fit for IPv6 because any 4 bits (half of a byte, or half of an octet) can be represented as a single hexadecimal digit. In other words, there are 16 unique combinations of 4 bits, and there are also 16 digits in a hexadecimal number system, so it is a perfect match. Because one hexadecimal digit can represent 4 bits, two hexadecimal digits can represent a single byte, or octet. For this reason, hexadecimal is commonly used in computer science, computer networking, and other areas of computer technology.

Note 4 bits is half a byte or half an octet, also known as a *nibble*. You will sometimes see alternative spellings such as *nybble* and *nyble*.

IPv6 Address Types

In Chapter 4 we will discuss the hexadecimal format of an IPv6 address and examine in detail all the different IPv6 address types, including the ones we will discuss here. For now, you just need a brief introduction to four common address types that are used in many of the protocols discussed throughout this book.

Global Unicast Address (GUA)

An IPv6 global unicast address (GUA) is a globally unique and routable IPv6 address. It is equivalent to a public IPv4 address. A GUA begins with either a hexadecimal 2 or 3. A GUA can be either a source or destination IPv6 address.

The following is an example of a global unicast address:

```
2001:db8:cafe:1::100
```

Don't worry about the format of this address for now. We will explain it in Chapter 4.

Link-Local Unicast Address

A link-local address is a unicast address that is local only on that link. The term *link* refers to a logical network segment or subnet. Link-local addresses are limited to the particular link and are not routable beyond the local subnet.

An IPv6 device doesn't have to have a global unicast address but it must have a link-local address. In other words, any device that is IPv6-enabled requires a link-local address. The device must be capable of self-generating the link-local address on the IPv6-enabled interface.

Link-local addresses commonly begin with fe80, as in this example:

```
fe80::a299:9bff:fe18:50d1
```

Link-local addresses are typically created automatically by the host operating system, which is why you see these addresses already configured on devices with

Windows, Mac OS, and Linux operating systems. A link-local address can be either a source or destination IPv6 address.

Figure 2-1 illustrates the difference between an IPv6 global unicast address and a link-local address.

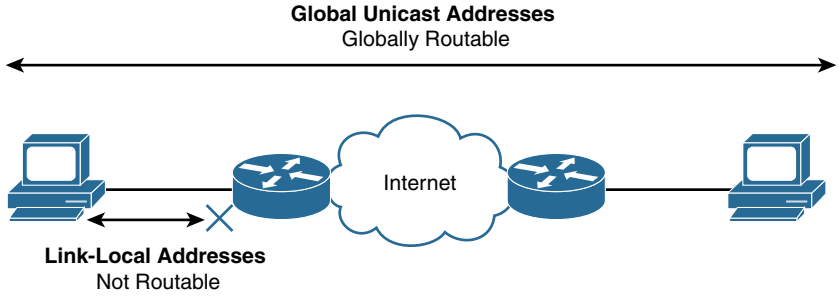

Figure 2-1 *Comparison Between IPv6 Global Unicast Address Addresses and Link-Local Addresses*

Unspecified Address

An IPv6 unspecified address is an all-0s address that indicates either the absence or anonymity of an IPv6 address. Unspecified addresses are used only as source addresses and never forwarded by an IPv6 router.

Solicited-Node Multicast Address

A solicited-node multicast address is a special multicast address that is used by some IPv6 protocols in lieu of the broadcast address that is used in IPv4. When an IPv4 broadcast address is encapsulated in an Ethernet frame, Ethernet uses a broadcast address for the destination MAC address. All devices on the Ethernet network (IPv4 subnet) will receive and process the frame. This means the Ethernet NIC must accept the frame and pass it up to the appropriate Layer 2 or Layer 3 protocol for processing. Some protocols, such as ARP, are not encapsulated in an IPv4 packet but the message may still be sent using an Ethernet broadcast address.

Note IPv6 does not have a broadcast address but does include an all-IPv6 devices multicast address. Chapter 4 discusses multicast address types.

This can be an inefficient process at times, such as in an ARP Request (see Figure 2-2). The ARP Request from PC1 is only trying to communicate with PC3. PC1 knows the IPv4 address of PC3 but not its Ethernet MAC address. However, the ARP Request is